AIME-Con 2026

Artificial Intelligence in Measurement and Education Conference (AIME-Con)

Volume 2: Works in Progress

October 5-7, 2026

The AIME-Con organizers gratefully acknowledge the support from the following sponsors.

Platinum

duolingo english test

Pearson

Gold

Curriculum Associates

i-Ready

*ets research institute

Silver

AAMC

Cambium Assessment

CollegeBoard

NBME

Supporter

edCount LLC because all students count

Eedi

MI MEASUREMENT INCORPORATED

Riverside Insights


University of
Pittsburgh
Learning Research and
Development Center

WEST COAST
ANALYTICS
complexity...simplified

WestEd
YEARS

duolingo english test

The future of language assessment is here

The Duolingo English Test is a computer adaptive test powered human-in-the-loop AI and supported by rigorous validity research. The test measures speaking, writing, reading, and listening skills, providing a deeper insight into English proficiency.



Built on the latest language assessment science

- Accessible by design, supporting test takers wherever they are for just $70
- Built on rigorous research and industry- leading security
- Integrates the latest assessment science and AI for accurate results
- Accepted by over 6,500 programs worldwide



englishtest.duolingo.com



Measure what matters

One connected learning story

Navvy+ brings skills-based MTSS data and standards-level classroom insight together in one platform, helping educators see where students are, what they need, and how to respond with greater confidence.

By connecting screening, progress monitoring, and classroom formative data, Navvy+ makes it easier to spot patterns, personalize support, and follow growth over time, without jumping between systems or piecing together disconnected reports.

Visit PearsonAssessments.com/navvy to learn more

800-627-7271 | PearsonAssessments.com
© 2026 Pearson Education, Inc. or its affiliates. All rights reserved. Pearson is a trademark, in the US and/or other countries, of Pearson plc

*ets
research
institute

AI research that empowers every learner

ETS Research Institute advances responsible AI research in learning, education and assessment.

We're applying decades of measurement science expertise to ensure AI-powered learning and measurement solutions are valid, equitable, fair and transparent. ETS Research Institute promotes the responsible use of AI by conducting rigorous research, creating practical tools and setting meaningful standards.

Learn more about our research on AI for assessment capabilities



Discover our AI research expertise at NCME AIME 2026:

- Validating AI-simulated students to draw conclusions about teacher practice
- Testing AI-generated reply consistency across model changes
- Promoting individual-level fairness in automated scoring — and reporting the accuracy tradeoff
- Verifying that AI-personalized math problems remain mathematically equivalent to standard items
- Documenting how AI-generated reading items can match statistics yet require different inferences

Connecting **Data** to **Instruction**



i-Ready Inform™

LEARN MORE

NBME

ADVANCING ASSESSMENT,
SUPPORTING OPTIMAL CARE

Through research and collaboration, NBME is evolving how we evaluate and support learners, with a focus on applying new technology to develop assessments that measure and build the knowledge and skills needed to provide optimal, effective care to all.

innovationsinassessment.org



AAMC

The AAMC Strengthens the World's Most Advanced Medical Care

Representing institutions that deliver the world's most advanced medical care, the AAMC provides programs and services that support the entire spectrum of education, research, and health care. We believe in medical education that prepares physicians and scientists to meet the nation's evolving health needs, support a diverse and culturally competent health care workforce, lead medical advancements that prevent disease and alleviate suffering, and more.

Learn more at aamc.org.

26-050 (01/26)

AAMC.ORG



Driving **impactful assessment** change in education through reliable and data-driven insights.

Advancing the Standard
View our Research

CambiumAssessment


CollegeBoard
College Board Is a Proud Sponsor of AIME-Con
Join our engaging sessions to learn how we're advancing innovative and responsible use of AI in educational measurement.
AI
ncme

©2026 National Council on Measurement in Education (NCME)

Order copies of this and other National Council on Measurement in Education (NCME) proceedings from:

National Council on Measurement in Education (NCME)
520 S. Walnut St. Box 2388
Bloomington, IN 47402
USA
Tel: +1-812-245-8096
ncme@ncme.org

ISBN 979-8-9983004-1-7

Preface



Introduction

The second NCME-sponsored Artificial Intelligence in Measurement and Education Conference (AIME-Con) brought together an interdisciplinary community of experts working at the intersection of artificial intelligence (AI), educational measurement, assessment, natural language processing, learning analytics, and technological development. As AI continues to transform education and assessment practices, this conference provided a critical platform for fostering cross-disciplinary dialogue, sharing cutting-edge research, and exploring the technical, ethical, and practical implications of AI-driven innovations in measurement and education.

Conference Theme - Measurement Science in AI-Integrated Assessment and Pedagogy

The second NCME-sponsored Artificial Intelligence in Measurement and Education Conference (AIME-Con) focused on illustrating ways that artificial intelligence is reshaping assessment, feedback, and pedagogy, while underscoring that these innovations must be anchored in rigorous measurement science. As AI systems expand the possibilities for scoring, item generation, adaptive tasks, and personalized instructional support, they simultaneously create new demands for evidence about what these systems measure, how well they measure it, and how they affect learning. The conference theme "Measurement Science in AI-Integrated Assessment and Pedagogy" frames psychometric principles—validity, reliability, fairness, and instructional effectiveness—as the foundation for responsible AI design and evaluation. It highlights the necessity of integrating measurement theory, data science, and learning research to guide the development and use of AI tools. The theme signals a commitment to advancing not only the capabilities of AI but also the evidentiary standards required to ensure these technologies enhance education.

The National Council on Measurement in Education

The National Council on Measurement in Education is a community of measurement scientists and practitioners who work together to advance theory and applications of educational measurement to benefit society. A professional organization for individuals involved in assessment, evaluation, testing, and other aspects of educational measurement, our members are involved in the construction and use of standardized tests; new forms of assessment, including performance-based assessment; program design; and program evaluation. Learn more about NCME, including our goals and our leadership, at www.ncme.org. We are grateful to the NCME.

NCME Special Interest Group on Artificial Intelligence in Measurement and Education

The AIME SIGIMIE seeks to advance the theoretical and applied research into AI of educational measurement by bringing together data scientists, psychometricians, education researchers, and other interested stakeholders. The SIGIMIE will discuss current practices in using Generative AI, approaches to evaluate their precision, accuracy, and areas where more foundational research is required into the way we test and measure educational outcomes. This group seeks to create a strong professional identity and intellectual home for those interested in the use of AI in many areas, including automated scoring, item evaluation, validity studies, formative feedback, and generative AI for automated item generation.

Proposal Requirements and Review Process for Work-in-Progress Papers

AIME-Con invited submissions of Working Papers, which were submissions of up to 1,000 words in the form of a structured summary. This format was designed for work-in-progress or pilot studies. Working Papers required a title, short abstract, and followed a structured format with the following headings:

- Background
- Aims
- Sample(s)
- Methods
- Results (or Anticipated Results)
- Conclusions (or Anticipated Contributions)
- References, tables, and figures included as needed

Submissions were blinded for peer review.

Submissions were evaluated by members of the review committee using a rubric that evaluated the following dimensions:

- **Relevance and community impact**: pertinence to the conference theme and AI in measurement and education community broadly, and potential contribution to current discussions and challenges in the field.
- **Scholarly or Practical Significance**: scholarly merit or practical importance of the work, and potential impact on theory, practice, or policy.
- **Methodological rigor**: coherence and appropriateness of the proposed methods, techniques, and approaches; and soundness of the overall research design.
- **Strength and trustworthiness of the results and conclusions**: whether the proposed analysis and interpretation methods yielded trustworthy results and conclusions that contribute valuable knowledge to the field.

For the purposes of this conference, "AI" was defined broadly to include rule-based methods, machine learning, natural language processing, and generative AI/large language models. Reviewers provided constructive feedback and overall recommendations to ensure that accepted sessions reflected both scholarly merit and practical value to the AI in measurement and education community.

Organizing Committee

NCME Leadership

Kadriye Ercikan, Ph.D. (President)
Susan Lyons, Ph.D. (Executive Director)

Conference Chairs

Joshua Wilson, Ph.D., University of Delaware
Christopher Ormerod, Ph.D., College Board
Magdalen Beiting-Parrish, Ph.D., EdAlfy, CUNY Graduate Center

Proceedings Chair

Nitin Madnani, Ph.D., Duolingo

Program Chairs

Joshua Wilson, Ph.D., University of Delaware
Christopher Ormerod, Ph.D., College Board
Magdalen Beiting-Parrish, Ph.D., EdAlfy, CUNY Graduate Center
Polina Harik, Ph.D., NBME

Program Committee

Program Chairs

Joshua Wilson, Ph.D., University of Delaware
Christopher Ormerod, Ph.D., College Board
Magdalen Beiting-Parrish, Ph.D., EdAIfy, CUNY Graduate Center
Polina Harik, Ph.D., NBME

Reviewers

Ketan , University of Massachusetts Amherst
Tazin Afrin, NBME
Norah Almousa, University of Pittsburgh
Talal Alzabidi, University of North Carolina Greensboro
Kylie Anglin, University of Connecticut
Olukayode Apata, Texas A&M University
Sergio Araneda, Caveon
Burcu Arslan, ETS Research Institute
Aydan Azimzade, Teachers College, Columbia University
Damilola Babalola
Justin Barber, Pearson
Mark Beck, Measurement Incorporated
Ahmed Bediwy, Pearson
Kashish Behl, ETS
Magdalen Beiting-Parrish, EdAIfy, CUNY Graduate Center
Ummugul Bezirhan, Boston College
Obed Boateng, Virginia Commonwealth University
Kelly Boden, Amira Learning
Brad Bolender, Finetune by Prometric
Laine Bradshaw, Pearson
Jill Burstein, Duolingo
Leon Camus, DIPF | Leibniz Institute for Research and Information in Education
Roti Chakraborty, American Institutes for Research
Shubham Chakraborty, Arizona State University
Jie Chen, Measurement Incorporated
Siyuan Chen, Duolingo Inc.
Ikkyu Choi, ETS
Jeongwon Choi, Vanderbilt University
S. Thomas Christie, Renaissance Philanthropy
Michael Chrzan, Center for Educational Data Science and Innovation, University of Maryland
Ruhan Circi, American Institutes for Research
Dante Cisterna, ATLAS - University of Kansas
Amy Clark, ATLAS: University of Kansas
Susan Davis-Becker, ACS Ventures LLC
Victoria Delaney, San Diego State University
Onur Demirkaya, Riverside Insights
Rahul Divekar, Bentley University
Svetlana Dmitrieva, University of South Carolina
Enis Dogan, NCES

Andrew Emerson, National Board of Medical Examiners
Farina Faiz, George Mason University
Fen Fan, National Board of Medical Examiners
Meng Fan, HUMRRO
Michael Fauss, ETS
Yetunde Fawehinmi, Texas A&M University
Peter Foltz, University of Colorado - Boulder
Julius Frost
Jie Gao, McGill University
Guher Gorgun, University of Georgia
Irina Grabovsky
Edith Graf, Proof-of-Concept Learning LLC
Joe Grochowalski, College Board
Yi Gui, Measurement Inc
Suhwa Han, Cambium Asessment
Polina Harik, NBME
Brian Harrold, University of South Carolina
Yong He
Michael Hemenway, Pearson
Langdon Holmes
Shahin Hossain
Jing Huang, Purdue University
Sijia Huang, Indiana University Bloomington
Wei Huang, University of Alabama, College of Education
Yue Huang, Measurement Incorporated
Fareya Ikram, University of Massachusetts Amherst
Megan Imundo, WGU Labs
Gregory Jacobs, Pearson
Paul Jewsbury, Duolingo
Yanlin Jiang, NCCPA
Ji Yoon Jung, Boston College
Julie Jung, Harvard University
Tzur Karelitz, National Institute for Testing & Evaluation (NITE)
Farzan KarimiMalekabadi, Edmentum
Euigyum Kim, Sogang University
Eileen Koh, Purdue University
Harvey Ngoe Kolle, University of Alberta
Miryeong Koo, University of Illinois at Urbana-Champaign
Tianqi Lang
Antonio Laverghetta Jr., Wesleyan University
Thuy Dieu Le, University of Massachusetts Amherst
Sophia Lebaschi, California State University Northridge
Haneul Lee, University of Massachusetts Amherst
Jimin Lee, Clark University
Will Lee, UMass Amherst
Youlim Lee, University of Massachusetts Lowell
Arun Balajiee Lekshmi Narayanan, University of Pittsburgh
Jennifer Lewis, University of Massachusetts Amherst
Jujia Li, University of Alabama
Tianwen Li, University of Pittsburgh
Zhifei Li

Guangming Ling, ETS
Helena Liu, West Lafayette Junior/Senior High School
Xiang Liu, ETS
Yong Luo, Pearson
Meng Lyu
Jing Ma, University of Iowa
Ye Ma, Amazon Web Services
Hotaka Maeda, Smarter Balanced
Henry Makinde
Michael McCarthy, Education Collaboratory at Yale University
Vandhana Mehta, Pearson
Steven Moore, George Mason University
Josh Moskowitz, Acuity Insights
William Muntean, National Council of State Boards of Nursing
NaRa Nam
Sungjin Nam, ACT Education Corp.
Anna Nasyrova, University of Massachusetts Amherst
Chunling Niu, The University of Texas at San Antonio
Teresa Ober, ETS
Burhan Ogut, American Institutes for Research
Christopher Ormerod, College Board
Corey Palermo, Measurement Incorporated
Jung Yeon Park, George Mason University
Hallie Parten, University of Virginia
Benjamin Pierce, Institute for Intelligent Systems
Prashreet Poudel, Western Carolina University
Jiten Pradhan, American College of Cardiology
Emily Proctor, College of Engineering, Purdue University
Aleena Raj, Kaplan North America
Ranjidha Rajan, MSU Denver
Brandi Robinson
Alex Rudniy, Drew University
Christopher Runyon, National Board of Medical Examiners
Ruchi Sachdeva, Pearson
Tahira Saddiqa, university of North Texas
Zachary Schultz, Cambium Learning Group Inc.
Jonathan Schweig, RAND
Qian Shen, University of Florida
David Shin, American Board of Psychiatry and Neurology
Jinnie Shin, University of Florida
Anthony Shiver, Law School Admission Council
Gozde Sirganci, Southern Methodist University
William Skorupski, Amira Learning
Dan Song, The University of Iowa
Angela Starrett, University of South Carolina
Gordon Stein, Renaissance Philanthropy
Sunday Stein, San Diego State University & University of California
Javier Suárez-Álvarez, University of Massachusetts Amherst
Susha Suresh, Acuity Insights
Judy Tang, Westat
Steven Tang, eMetric LLC

Ethan Tenison, Amira Learning
Danielle Thomas, Carnegie Mellon University
Yu Tian, Arizona State University
Lingbo Tong, University of Wisconsin–Madison
Kazi Md Rukan Uddin, University of Houston
Kirk Vanacore, Cornell University
Claudia Ventura, University of Connecticut
Yen Vo, The University of Iowa
Matthias von Davier, Boston College
Mai Vuong, University of Pennsylvania
Cole Walsh, Acuity Insights
Kuo Wang, Southern Methodist University
Ling Wang
Ting Wang
Zuowei Wang, Educational Testing Service
Field Watts, ETS
Yumou Wei, Carnegie Mellon University
David Whitcomb, ATLAS, University of Kansas
John Whitmer, Learning Data Insights
Joshua Wilson, University of Delaware
John Workman, EQAO
Lixin Wu, University of Illinois at Urbana-Champaign
Sirui Wu, University of British Columbia
Yi-Chen Wu, National Center on Educational Outcomes, University of Minnesota
Yi-Fang Wu, Cambium Assessment Inc.
Michael Xiao, University of Washington
Weixuan Xiao, Teachers College, Columbia University
Dawei Xie, Northwestern University
Tianyu Xu, Teachers College, Columbia University
Yue Xu, University of Florida
Kang Xue, Cambium Assessment
Rohit Yadav, University of Kansas
Daijin Yang, Northeastern University
Yaxuan Yang, University of Georgia
Yiting Yao, Florida State University
Farzana Yasmin, University of Alberta
Michelle Yin, Northwestern University
Jiseung Yoo
Jessica Younger, Digital Promise
Kefan Yu
Idris Yusuf, University of Houston
Hongyi Zhai, North Oconee High School
Liang Zhang, University of Michigan
Qiongyao Zhang, Robert Morris University
Shan Zhang, University of Florida
Yuxiao Zhang, Purdue University
Zhehan Zhang, Clemson University

Table of Contents

A Cognitive Lab Study of Student AI Chatbot Use in Online Coursework
Jinah Choi, Sonya Powers, Farzan Karimi-Malekabadi and Michelle Barrett 1

A Preliminary Semantic Screen for IRT Local Dependence
Qian Shen 9

Evaluating Multi-Phase and Granular Strategies for Evidence-Anchored Automated Scoring
Alessia Marigo, Laura Wright, Linda Malkin and Xin Xie 17

Consensus without Accuracy: Investigating LLM's Recovery of Item Difficulty Using Paired Comparisons
Michael Leon Chrzan and Benjamin Domingue 24

Implementing Generative AI Assessment Guidelines in a Mexican University
Melchor Sánchez-Mendiola, Enrique Buzo-Casanova, Elibidú Ortega-Sánchez, Manuel García-Minjares and Adriana Cruz Juárez 37

TRACE: Automated Multi-Granularity Analysis of Text Revision
Yu Tian, Katerina Christhilf, Andrew Potter, Jessica Early, Steve Graham and Danielle McNamara 45

Comparing θ Representation Strategies for Reconstructing Item Characteristic Curves
Sungjin Nam 52

Assessing Large Language Model Performance in Post-Certification-Examination Comment Categorization
Huaping Sun, Kristin O'Brien, Colleen Burke Kave, David Shin, Qiao Lin and Jeffrey Marc Lyness 57

Rethinking Pilot Data: Evaluating LLM Synthetic Data for Scale Development
Margarita Olivera-Aguilar, Samuel H. Rikoon, Paul D. Bailey, Michael B. Kruse and Kamal Middlebrook 62

Hierarchical Attention Network Architecture for Automated Response Time Prediction
Wei Shuang Schneider 69

Beyond Agreement: Calibrating Automated Scoring to Human Judgment Using Control Scripts
Mark Dulhunty, Alex Codoreanu and Nathan Zoanetti 79

Human-in-the-Loop Gemini Item Generation for Large-Scale MCQ Banks
Melchor Sánchez-Mendiola, Milton García-Lima, Elibidú Ortega-Sánchez, Manuel García-Minjares and Enrique Buzo-Casanova 86

Score Comparability Limits with Common AES Validation Measures
Tim Moses, Seongeun Kim and Nick Trout 93

Finding Evidence of Communication Behaviors using LLMs
Tazin Afrin, Saed Rezayi, Janet Mee, Polina Harik and Le An Ha 100

Designing and Evaluating a Measurement-Oriented RAG Framework for ELL Writing Assessment
Daeryong Seo 104

Automated Item Evaluation: Predicting Item Acceptance and Rejection using LLM-Generated Critiques
Hotaka Maeda and Yikai Lu 113

Production-Ready Automated Item Generation in Educational Assessment: Integration with Operational Workflows
Hotaka Maeda, Kargi Chauhan and Tharunya Chandrashekar 130

Using Generative AI to Simulate Item Responses by Skills Insight Score Band
Jianshen Chen, Chansoon Lee, Kylie Gorney and Judit Antal 140

Exploring AI-driven Methods for Pre-calibrating Difficulty of Mathematics Items with Visual Content
Leah Walker and Jinnie Choi 145

Evaluating Multiple Models for Predicting Item Difficulty in a Principled Assessment Design
Alexandra Lane Perez, Christina Schneider, Sangdon Lim, Garron Gianopulos and Kang Xue 152

Transformer-Aided Detection of Gaming in Constructed-Response English Language Assessments
Melanie Sharif, Scott Hellman, Martha Bellows and Sue Lottridge 159

Calibrated, Interpretable Automated Scoring with LLM Likelihoods
Thomas Christie, Markus Hauru and Anna Rafferty 168

Predicting Item-to-Range Performance Level Descriptor Matches with Structured Language Model Reasoning
Kang Xue, M. Christina Schneider, Sangdon Lim, Garron Gianopulos and Alexandra Perez . 177

Using LLM Judges' Paired Comparisons to Estimate Mathematics Item Difficulty
Lanrong Li, Mohammad A. A. Abulela and Matthew Gushta 184

Math-Specialized or General-Purpose? A Comparison of MathBERT and DeBERTa-v3-large for Automated Scoring of Mathematics-Explanation Items
Gregory M. Jacobs, Ahmed H. Bediwy and Martha Bellows 193

How Much Does Hyperparameter Tuning Actually Help? An Efficiency Survey for Fine-Tuning Transformers to Score Mathematics-Explanation Items
Gregory M. Jacobs, Ahmed H. Bediwy and Martha Bellows 200

Responsible use of generative AI when creating reading comprehension questions: Inference matters
Zuowei Wang, Michael Flor, Jiyun Zu, Tenaha O'Reilly and Wanjing Anya Ma 207

Fine-tuning DeBERTa-v3 to Automate Spatial Language Classification
Sam Agnoli, Qingzhou Shi and David Uttal 211

CDM Q-Matrix Discovery with LLMs: Fusing Domain Knowledge with Empirical Evidence
Susu Zhang and V. N. Vimal Rao 216

Composite Scores vs. Preference Rankings: Measuring Architecture Effects in LLM Feedback
Harvey Ngoe Kolle, Carrie Demmans Epp, Amna Liaqat and Maria Cutumisu 231

Identifying Communication Profiles from Physician–Patient Conversations Using Large Language Models
Reyhaneh Hosseinpourkhoshkbari, Yash Bipin Jain and Richard Golden 235

LLM-Based Pairwise Judgment for Math Item Parameter Modeling Using Workflows and Agents
Suhwa Han and Frank Rijmen 248

Beyond Volume: How Cognitive Fingerprints Predict Residual Gain in AI Tutoring
Aleena K Raj and Nina Deng 260

Automatic Prompt Engineering for Generative AI–Based Essay Scoring
Yue Huang 267

A Comprehensive Evaluation of GenAI-based Items in a Medical Examination
Yanlin Jiang, Marcus Walker, Andrew Dallas, Aquia Richburg, Nikole Gregg and Brittany Corrigan 275

When One Benchmark Hides Many Truths: Mixture-IRT and LLM Difficulty Prediction
Enis Dogan and Burhan Ogut 282

LLM-As-An-Assessor: Can Open-Weight LLMs Assess Computational Thinking via Student-Designed Embodied Games?
William Lee, Sai Gattupalli, Ivon Arroyo and Brendan O'Connor 289

Funnel Plot Analysis of Examinee Performance in Longitudinal Assessment
Aquia Richburg and Marcus Walker 298

How Much Training Data Is Enough? Fine-Tuning LLMs for Short-Answer Scoring
Joshua A McGrane and David Torres Irribarra 306

Evaluating Machine Learning, Transformer, and LLM Approaches for Large-Scale Automated Essay Scoring
Cyrus Sai-Cheong Chan, Charles Anifowose, Su Zhang, Winnie Wing-Yee Tse, Nizam Radwan and Hyunah Kim 315

Construct Validity of Small-Sample Transformer Scoring Models: A Mechanistic Interpretability Approach
Michael P. Hemenway and Martha Bellows 322

Applying Evidence-Centered Design to Automated Evals of AI-Powered Assessment Systems
Kristen DiCerbo, Britte Haugan Cheng and John Whitmer 329

MIST: Misconception Identification and Sycophancy Tracking
Lionel Meng, Lingbo Tong and Icy Zhang 335

Building Teacher Capacity to Generate and Evaluate Math Items with Small LLMs
Minseok Kim 344

Principled Approaches to Building AI Partners for Assessing and Supporting Small-Group Collaboration
Peter W Foltz, Chelsea Chandler and Mon-Lin Monica Ko 350

Measuring Writing Proficiency Through Keystroke Process Data: A Cross-School Generalizability Study
Damilola Babalola, Naga Buddarapu, Caleb Scott, Piotr Mitros, Paul Deane and Collin Lynch 358

Self-Revising Agents as Item Writers: Benchmarking Against Professional Item Writing
Steven Tang, Zhen Li and Richard Patz 369

Evaluating Neural Network-Based IRT in Multistage Adaptive Testing with Small Item Banks
Seong Eun Hong 377

Validating LLM-Rated Item Features for Explanatory Item Response Models
Ahmed H. Bediwy and Mubarak O. Mojoyinola 382

Efficacy of Student–AI Co-Authored Math Word Problems in an Intelligent Tutoring System
Kole Norberg, April Murphy, Steve Fancsali and Steve Ritter 389

Assessing Small Language Models as Decimal-Arithmetic Tutors: A Measurement Framework
Mai Que Vuong, Shahana Ahmadli, Michelle Zhou, Minseok Kim, Shruti Mehta, Talita de Paula Cypriano de Souza and Seiji Isotani 397

LLM as Investigator and Judge in Pairwise Comparisons for Item Parameter Modeling
Brian Harrold, Vincent Fakiyesi, Maria D'Brot, Susan Lottridge, Justin Barber and Michael Hemenway 405

LLM-Based Metadiscourse Analysis of Demographic Variations in ELL Hedge Use
Xinyu Xu and Tianyu Xu 416

Beyond the Grid: Auditing Code-Editing Search for Automated Essay Scoring
Julius Frost 423

LLM Difficulty Prediction Across Three L2 Task Types and Two Examination Forms
Meng Lyu, Moses Oluoke Omopekunola and Feng Wang 441

Measuring Science Teachers' Generative AI Literacy for Classroom Assessment: Performance Assessment
Ruiping Huang and Yue Yin 448

Do Synthetic Students Thrive? Validity and Fairness of Generated Assessment Data
Neba Nfonsang and Temple Lovelace 459

AI-Enabled Quality Assurance for Multiple-Choice Assessment Items
Steven James Moore and Nicholas Diana 464

More than Words: Modeling What Teachers Praise and How They Praise It
Jiseung Yoo, Paiheng Xu and Jing Liu 472

Does an Automated Source Integration Score Respond to Revision? Extending Evidence of Construct Validity with a Feedback Experiment
Andrew Potter, Yu Tian, Kaitlin Van Houghton, Manmeet Singh, Renu Balyan, Maria Goldshtein, Laura K. Allen and Danielle S. McNamara 483

A Cognitive Lab Study of Student AI Chatbot Use in Online Coursework

Jinah Choi, Sonya Powers, Farzan Karimi-Malekabadi, and Michelle D. Barrett
Edmentum

Abstract

This cognitive lab study examined how secondary students in Grades 7-12 interacted with an AI-enabled chatbot while completing online course activities. Using chat logs, facilitator observations, post-session reflections, and survey responses, we analyzed spontaneous and lightly seeded interactions to evaluate whether the chatbot supported its intended theory of action: helping students continue learning without simply completing work for them. Findings highlight the importance of typing fluency, answer-boundary handling, safety and privacy response patterns, and student-facing feedback mechanisms. This study illustrates how cognitive lab evidence can inform ongoing validation.

1 Introduction

AI-enabled chatbots are increasingly being explored as support tools in online learning environments. Prior reviews suggest that educational chatbots can support learning by providing immediate assistance, feedback, and opportunities for interaction, while also emphasizing that effectiveness depends on design, implementation context, and the quality of student–system interaction (Kuhail et al., 2023). Recent work on generative AI in K–12 education similarly highlights both the promise and unresolved risks of AI-supported instruction, including concerns related to accuracy, overreliance, privacy, equity, and appropriate integration into learning activities (Alfarwan, 2025). A recent meta-analysis further suggests that chatbots may have a small-to-moderate positive effect on learning performance, while cautioning against treating chatbot use as generically beneficial across educational contexts (Laun et al., 2025).

Despite this growing body of work, there remains limited direct evidence on how K–12 students actually interact with AI tutors embedded in online coursework. Much of the existing literature has focused on chatbot design, empirical effectiveness, limitations, or student perceptions, but less is known about how learners use these tools while completing real course activities. This gap is important because students may use AI tutors in ways designers do not fully anticipate: they may ask vague or informal questions, use incomplete language, seek direct answers, challenge the chatbot, disclose sensitive information, or ask for help with tasks outside the chatbot's intended role.

The present study is motivated by the need to understand these student-AI interactions before AI learning assistants are introduced to a broader population of learners. A cognitive-lab approach is useful because it allows researchers to observe how students use a chatbot in context: when they decide to ask for help, how they formulate questions in their own words, what kinds of informal or incomplete language they use, whether they ask follow-up questions, how they respond when support is helpful or unclear, and whether they attempt to challenge the system or use it in unintended ways. These interaction patterns are difficult to anticipate through expert review alone, but they are central to evaluating whether an AI tutor is likely to function as intended with real students. Within this broader evaluation, answer-boundary behavior is especially important because the chatbot must provide useful guidance without giving copy-ready answers or reducing the cognitive work students need to do. By examining both authentic student use and high-priority interaction situations in a structured setting, this study illustrates how cognitive-lab evidence can

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 1–8
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

inform early evaluation of AI learning assistants before broader implementation.

2 Research Questions

This study is guided by the following research questions.

1. **How do secondary students use an AI-enabled chatbot during independent online coursework?**
   This question examines students' authentic use of the chatbot, including when they decide to ask for help, what kinds of questions they ask, how they formulate requests in their own words, and whether they continue or stop interacting after receiving a response.
2. **How does the chatbot support, redirect, or extend student thinking across naturalistic and lightly seeded interactions?**
   This question focuses on the pedagogical quality of chatbot responses, including whether the chatbot clarifies concepts, helps students reason through problems, prompts students to explain their thinking, supports checking strategies, addresses misunderstandings, and encourages students to continue the learning task.
3. **What do high-priority interaction episodes reveal about opportunities to strengthen pedagogical scaffolding?**
   This question examines interaction moments that are especially important for AI tutor evaluation, including vague help requests, answer checking, requests for more direct support, off-task use, safety or privacy boundaries, and situations where the chatbot must balance helpfulness with preserving student cognitive engagement.
4. **How do students describe their experience using the chatbot, and what design and evaluation considerations emerge from their observed and reported interactions?**
   This question examines student-reported evidence from post-session reflections and surveys, including perceived helpfulness, ease of use, response clarity, productive struggle, trust, feedback-feature interpretation, and likely future use. It also considers broader implications for evaluating AI tutors with the intended student population.

3 Method

3.1 Study Design

This study used a one-on-one cognitive-lab design to examine how secondary students interacted with an AI-enabled chatbot while working independently on an online course. The design focused on authentic student–chatbot interaction: when students chose to ask for help, how they formulated requests in their own words, how they responded to chatbot guidance, and whether they continued, redirected, or stopped the interaction.

The session prioritized naturalistic use before any prompted scenarios were introduced. This allowed the study to capture students' spontaneous help-seeking behavior during coursework while also providing an opportunity, later in the session, to observe high-priority interaction types that may not occur naturally in a short session, such as answer-seeking, system-challenging behavior, off-task use, safety or privacy boundaries, and course-related support questions.

3.2 Participants and Setting

Participants were 13 secondary students completing assigned online course activities across English/ELA, mathematics/algebra, science/biology, and social studies/U.S. history. The completed sample included students from Grades 7–12. Because the study was qualitative and exploratory, grade level, content area, prior platform experience, and other student characteristics were treated as contextual variation rather than as variables for formal comparison.

Each session was conducted remotely with a facilitator. Students worked through assigned course content aligned to their grade level or subject area and were reminded that the session was not a test of their academic knowledge. The emphasis was on observing how students used the chatbot during coursework, particularly when they were confused, stuck, wanted a hint, wanted to check their thinking, or needed help deciding what to do next.

3.3 Session Procedure

Each session lasted approximately 45 minutes and included orientation, naturalistic chatbot use,

lightly seeded scenarios, post-session reflection, and a short survey.

During the orientation, the facilitator explained the purpose of the session, reminded students that they were not being graded, and encouraged them to use the chatbot naturally. Students were told that the study focused on what they chose to ask and how the conversation unfolded, rather than whether they answered course questions correctly.

During the naturalistic-use phase, students worked independently through assigned online course content for approximately 18–20 minutes. They could use the chatbot whenever it felt natural or helpful, including to ask for explanations, request hints, check their thinking, ask what to do next, or try to get the answer. The facilitator avoided giving specific scenarios during this phase unless the student became stuck or ran out of things to ask.

During the lightly seeded scenario phase, the facilitator introduced a small number of brief "what if" situations when needed. These scenarios were designed to broaden the range of observed interactions while still allowing students to compose their own chatbot messages. Scenario types included conceptual confusion, procedural uncertainty, direct or indirect answer-seeking, requests for the chatbot to complete work, vague or low-effort prompts, off-task questions, system challenges, safety or privacy boundaries, and course-navigation or support questions.

At the end of the session, the facilitator asked students a brief set of open-ended reflection questions about their intent in using the chatbot, perceived helpfulness, ease of asking for help, unexpected responses, and likely future use. These questions were intended to clarify students' experiences in their own words. Students then completed a short structured survey covering related constructs, including ease of use, response clarity, learning support, productive struggle, course-navigation support, and likely use.

3.4 Data Sources

The analysis drew on four sources of evidence: chatbot interaction logs, facilitator observation notes, post-session reflection responses, and student survey responses. Chatbot logs served as the primary behavioral record, capturing the sequence of student and chatbot messages during both naturalistic use and seeded scenarios.

Facilitator notes documented contextual information not always visible in the transcripts, including hesitation, typing burden, visible confusion, use of course tools, and return to the learning task. Post-session reflections captured students' own explanations of what they were trying to do, how helpful the chatbot was, what felt unexpected, and when they would use a similar tool. Survey responses provided descriptive evidence of students' perceptions of ease of use, clarity, learning support, productive struggle, trust calibration, course-navigation support, and likely future use.

3.5 Scenario Framework

The seeded scenarios were organized around interaction categories relevant to evaluating AI learning assistants. Conceptual-struggler scenarios examined requests for explanations, simpler wording, examples, rewording, concept comparisons, and study strategies. Procedural-errant scenarios examined reasoning checks, step-by-step help, answer comparison, and readiness to respond.

Answer-seeking scenarios examined direct requests for answers, indirect attempts to narrow choices, requests for completed work, and continued pushing after a hint. These scenarios were included because answer-boundary behavior is central to the intended role of an AI tutor: the chatbot should support learning without simply completing the task for the student.

Additional scenarios examined vague or low-effort input, off-task requests, topic switches, low motivation, chatbot challenges, requests for intentionally wrong answers, safety or privacy boundaries, and course-navigation or support questions. These scenarios expanded the range of observable interaction patterns beyond what might occur during a short naturalistic session.

3.6 Qualitative Analysis

Chat logs and facilitator notes were analyzed using structured qualitative coding to identify observable student–chatbot interaction patterns. Naturalistic-use and seeded-scenario episodes were coded separately to distinguish spontaneous behavior from prompted high-priority situations. Coding was conducted by one primary researcher using an initial codebook developed from the research questions, protocol, and seeded scenario framework. Codes included help-seeking goal,

prompt style, follow-up behavior, answer-seeking, response to hints, off-task or low-effort use, safety/privacy boundaries, and course-navigation or support-seeking questions.

During transcript review, the codebook was refined to capture emergent patterns such as typing burden, difficulty expressing mathematical or visual information, student persistence, and answer-boundary ambiguity. Episodes were interpreted by considering the student request, chatbot response, and student follow-up together. Post-session reflections and survey responses were analyzed as student-reported evidence of perceived helpfulness, ease of use, expectations for chatbot behavior, trust, and likely future use. Because the study was exploratory and based on single coding, findings are reported as preliminary evidence of student–chatbot interaction rather than as generalizable estimates of effectiveness.

4 Findings

4.1 Students Used the Chatbot for Practical, Task-Driven Learning Support

Students generally used the chatbot as a practical support tool during coursework rather than as a conversational agent. During naturalistic use, students used the chatbot for definitions, explanations, hits, equation support, answer checking, and help understanding confusing questions or passages. Their prompts ranged from short phrases, such as asking for a term definition, to longer contextual questions about history, science, or algebra. Some students provided partial reasoning or one-word responses while working through a problem, while others engaged in longer multi-turn exchanges.

This pattern suggests that students were able to recognize moments when chatbot support might be useful during independent work. Students did not need fully polished or adult-like prompts to initiate support. Instead, they often used informal, incomplete, or conversational language, and the chatbot was generally able to respond in ways that kept the interaction moving. In post-session reflections, students commonly described using the chatbot to understand content, clarify what a question was asking, check how to solve something, define unfamiliar words, or get help with a first step. Most students who commented on helpfulness reported that the chatbot supported what they were trying to do.

At the same time, student use varied in the amount of detail they provided and in their willingness to continue the conversation. Some students asked one focused question and returned to the activity, while others engaged in extended back-and-forth problem solving. This variation is important because it suggests that chatbot use during coursework is not a single uniform behavior. Students may use the same tool as a definition source, hint provider, reasoning checker, step-by-step tutor, or shortcut attempt depending on the task and moment of need.

4.2 Chatbot Responses Were Generally Supportive, but Boundary Handling Varied Across Interaction Contexts

Across naturalistic-use interactions, the chatbot generally responded in ways consistent with a learning-support role. When students asked for definitions, clarification, hints, equation support, or help understanding course content, the chatbot typically provided explanations, guiding questions, or step-by-step prompts. These responses often gave students enough support to continue the conversation, revise their thinking, or return to the coursework after receiving targeted clarification.

In the high-priority interaction contexts, the chatbot often maintained its instructional role but showed more variation. For vague or low-effort academic prompts, it usually asked for more context or helped the student identify a next step. For off-task requests or attempts to move outside the learning activity, it tended to redirect the student back toward coursework. For safety or privacy prompts, it generally avoided encouraging personal disclosure and moved the interaction toward safer or more appropriate forms of support.

The most nuanced pattern appeared in answer-seeking interactions. The chatbot often avoided directly naming a final answer and instead offered hints, reasoning steps, or strategies for evaluating options. However, some responses still approached the answer boundary by eliminating unlikely choices, confirming narrow intermediate answers, identifying key elements in answer choices, or providing a response structure that made the final answer easier to infer. These cases suggest that evaluating an AI tutor requires attention not only to whether it explicitly gives an answer, but also to whether its guidance preserves enough cognitive work for the student.

4.3 High-Priority Interactions Revealed Opportunities to Strengthen Pedagogical Scaffolding

High-priority interactions showed that the key issue was not simply whether the chatbot gave an answer, but how it supported student thinking when students asked for confirmation, requested more direct help, tried to narrow answer choices, or challenged the chatbot's response. In many cases, the chatbot maintained a learning-support role by offering hints, explanations, reasoning steps, or redirection rather than completing the work for the student.

Several interactions also pointed to opportunities for stronger pedagogical scaffolding. For example, when students asked whether an answer or calculation was correct, the chatbot could more consistently ask students to explain their reasoning or show how to check the answer. When students focused on answer choices, the chatbot could shift attention to the underlying concept, evidence standard, or decision rule before asking students to apply that reasoning themselves.

These cases are best interpreted as pedagogical decision points rather than isolated failures. The important question is whether the chatbot's response helps students reason through the task or takes over too much of that reasoning. This finding suggests that AI tutor evaluation should examine not only accuracy and answer boundaries, but also whether the tutor preserves student reasoning, productive struggle, and cognitive engagement.

4.4 Students Reported Positive Experiences, with Important Nuance Around Fact Checking, Effort, and Future Use

Student-reported evidence suggested that students generally viewed the chatbot as useful and easy to engage with. In post-session reflections, students described using the chatbot to clarify confusing content, understand questions, define unfamiliar terms, check reasoning, and get help when no teacher or other support person was immediately available. Survey responses were also positive, especially for ease of asking, response clarity, learning support, productive struggle, and likely future use.

However, students' positive experience depended partly on how easily they could express their questions through text. Some students found it straightforward to ask questions, but others needed time to figure out what to ask or how much context to provide. This was especially relevant for tasks involving graphs, tables, mathematical notation, visuals, or longer passages. These observations suggest that student experience should be evaluated not only by whether students liked the chatbot, but also by how much effort was required to make the chatbot useful.

Students generally found the chatbot understandable and helpful, but survey patterns suggested that students may value chatbot support while still needing guidance on when to rely on it, when to check another source, and how to respond when an answer seems incomplete or confusing. Overall, students saw plausible future value in the chatbot for independent coursework, but their reported experience points to the need to take into consideration differing levels of AI literacy when considering student evaluations of AI usefulness.

5 Discussion and Significance

This study illustrates the value of cognitive-lab evidence for evaluating AI tutors as interactional learning tools in K–12 online learning environments. As AI becomes embedded directly within coursework platforms, the role of the chatbot is no longer limited to answering isolated student questions. It may function as a clarifier, hint provider, reasoning partner, checker, and feedback interpreter when students work independently and a teacher is not immediately available. Evaluating such tools therefore requires attention not only to response accuracy, but also to how students seek help, how the tutor responds pedagogically, and whether the interaction supports continued student reasoning.

A central implication is that effective AI tutoring depends on pedagogical moves, not only on guardrails. When students ask for confirmation, more direct help, or assistance choosing among options, the strongest response is not necessarily to refuse help or to provide the answer. More productive scaffolding may involve asking students to explain their reasoning, modeling how to check an answer, focusing attention on the underlying concept or evidence standard, and returning responsibility for the final decision to the student. This reframes answer-boundary concerns as a broader instructional design question: how can an AI tutor provide enough support to keep students moving while still preserving productive struggle, cognitive engagement, and student ownership of the work?

The findings also raise important questions about text-mediated help seeking in the current generation of AI-supported learning tools. Students must first represent their confusion, partial understanding, or task context in language before the chatbot can respond productively. In this study, students varied in how much context they provided, how they phrased uncertainty, and how easily they translated mathematical, visual, or passage-based information into a prompt. These observations should not be interpreted as generalizable differences among students, but they point to a broader issue for K–12 AI tutor design: as chatbot roles expand, researchers and developers need to consider how students communicate with AI systems, not only what those systems return. They also need to provide support to students with varying levels of AI literacy by providing resources up front to help students engage successfully with these new tools. Just because students are situated in a world where technology and AI tools are becoming pervasive throughout educational and non-educational contexts does not mean that all of the students are equally equipped to make use of the tools. AI chatbots can provide a useful tool to all students, but only if there is a level playing field of AI literacy that the students begin with, and this assumption cannot be made from the outset.

Moreover, typing skills were found to play an important role in the level of effort students had in engaging with the chatbot. Students who are growing up with autocomplete, predictive text, speech-to-text tools, and generative AI may not interact with text-only chat interfaces in the same way as earlier cohorts of internet users. For some students, typing a long, precise prompt may be a natural part of help seeking; for others, it may create friction before the tutoring interaction even begins. This does not mean that text-based chat is inappropriate for learning, but it does suggest that future AI tutors may need multiple ways for students to express what they need, especially for visual, mathematical, graph-based, or passage-heavy tasks.

Methodologically, the study shows why naturalistic use and lightly seeded scenarios are useful in combination. Naturalistic use captured how students chose to use the chatbot during a coursework activity. Seeded scenarios added evidence about less frequent but instructionally important situations, including answer-seeking, low-effort prompts, off-task use, and safety or privacy boundaries. These scenarios are especially important in K–12 contexts because rare events can still carry substantial instructional or safety implications. A short cognitive lab cannot estimate how often such events will occur, but it can identify the kinds of interactions that should be monitored, refined, or studied more systematically before broader implementation.

The significance of this work is therefore not limited to the specific chatbot examined here. More broadly, the study offers a framework for including qualitative research with target users of AI-enabled educational programs into larger evaluations studies. For AI chatbots and tutoring systems in particular, such evaluation should ask whether students can seek help in their own words, whether the tutor scaffolds rather than replaces student thinking, whether it handles boundary conditions in instructionally appropriate ways, and whether students perceive the support as useful, understandable, and worth using again. For K–12 AI learning systems, the key question is not simply whether a chatbot can generate a correct response, but whether it can participate in a learning interaction that helps students think, persist, and remain responsible for their own work.

References

AA Abdullah Alfarwan. 2025. Generative AI use in K–12 education: A systematic review. *Frontiers in Education*, 10:1647573.

Mohammad Amin Kuhail, Nazik Alturki, Salwa Alramlawi, and Kholood Alhejori. 2023. Interacting with educational chatbots: A systematic review. *Education and Information Technologies*, 28:973–1018.

Martin Laun and Fabian Wolff. 2025. Chatbots in education: Hype or help? A meta-analysis. *Learning and Individual Differences*, 119:102646.

Fengchun Miao and Wayne Holmes. 2023. *Guidance for Generative AI in Education and Research*. UNESCO, Paris.

Griffin Pitts, Viktoria Marcus, and Sanaz Motamedi. 2025. Student perspectives on the benefits and risks of AI in education. In *Proceedings of the 2025 ASEE Annual Conference & Exposition*, Montreal, Quebec, Canada.

Appendices

Appendix A. Student Cognitive Lab Session Flow

Segment	Approx. Time	Purpose
Welcome and orientation	3-5 min	Explain purpose, reduce performance pressure, introduce screen sharing
Integrated coursework and chatbot use	~20 min	Observe natural student use while working independently
Seeded scenarios	8-10 min	Elicit high-priority interaction types not observed naturally
Post-session reflection	3-5 min	Clarify student intent and perceived use context
Short survey and closing	3-5 min	Collect brief contextual perception data and close session

Appendix B. Seeded Scenario Framework

Scenario Category	Purpose
Conceptual Struggler	Observe how students seek help when they are confused about a concept, term, question, or task demand.
Procedural Errant	Observe how students seek support when they have some understanding but are uncertain about their process, reasoning, or answer choice.
Answer Seeker	Observe how students attempt to get closer to the answer or shortcut the task.
Unengaged Learner	Observe how the chatbot responds when students use low-effort, vague, off-task, or avoidance-oriented language.
System-Breaking / Challenge	Observe how students and the chatbot respond when the student tests the system or challenges the chatbot's response.
Crisis / Harm / Safety Boundary	Observe how the chatbot handles general or fictional safety, distress, or personal-disclosure boundary situations.
Course Navigation / Logistics	Observe whether students use the chatbot for course-level or platform-support questions, not only academic content help.
Cross-Cutting Modifier: Bilingual or Multilingual Input	Allow students to use English, another language, or a mix of languages if that reflects how they would naturally ask for help.

Appendix C. Participant Characteristics

Participant Characteristic	Summary
Total completed participants	13
Student age range	12–18 years
Grade levels represented	Grades 7–12
Student gender	Female: n = 7; Male: n = 6
School type	Public school: n = 13
Prior use of the online learning platform	No prior use: n = 10; Yes, a few times: n = 2; Not sure: n = 1
School location context	Suburban: n = 7; Small town: n = 4; Rural: n = 1; Urban: n = 1
School size	Fewer than 500 students: n = 2; 500–999 students: n = 5; 1,000–1,999 students: n = 4; 2,000 or more students: n = 2
Formal learning accommodation plan reported	Yes: n = 2; No: n = 11
Language background	English as primary language: n = 12; Bilingual or multilingual: n = 1
Geographic representation	8 U.S. states plus the United Kingdom
Assigned content areas	English/ELA: n = 4; Math/Algebra: n = 4; Science/Biology: n = 2; Social Studies/U.S. History: n = 3

Appendix D. Post-Session Reflection Questions

Question	Main Focus	Student-Facing Question and Probes
Q1	Initial chatbot-use intent	**When you first used the chatbot today, what were you trying to do?** For example, were you trying to understand something, get help with a question, check your answer, figure out what to do next, or just see what the chatbot could do? **Follow-up: Did the chatbot help you with what you were trying to do, or not really?**
Q2	Ease of asking for help	**Was it easy or hard to figure out how to ask the chatbot for help today?** For example, did you know what to type, or did it take a little time to figure out? Did you ever have to ask your question a different way?
Q3	Unexpected chatbot responses	**Was there any time when the chatbot answered in a way you did not expect?** What was different from what you expected? What did you do next?
Q4	Likely use context	**Imagine you were working on an online course by yourself after school or at home, and a teacher was not available right away. When do you think you might use a chatbot like this?** For example, would you use it when you are stuck, when you want a hint, when you want to check your answer, when you are not sure what to do next, or for something else?

Appendix E. Student Survey Items

(1:Strongly disagree; 2:Disagree; 3:Neither agree nor disagree; 4:Agree; 5:Strongly agree)

#	Construct	Item
1	Ease of asking	I could ask the chatbot questions in my own words.
2	Ease of asking	I knew how to ask the chatbot for help when I needed it.
3	Clarity	The chatbot's responses were easy to understand.
4	Clarity	The chatbot explained things in a way that made sense to me.
5	Learning support	The chatbot helped me think through the activity.
6	Learning support	The chatbot helped me keep working when I was unsure what to do.
7	Productive struggle tolerance	It was okay when the chatbot gave me a hint instead of giving me the answer right away.
8	Productive struggle tolerance	I was willing to keep trying when the chatbot guided me instead of just giving the answer.
9	AI Literacy	I knew that I should check important answers against the lesson, course feedback, or teacher guidance.
10	AI Literacy	The chatbot always provided correct information.
11	AI Literacy	If the chatbot gave an answer that seemed confusing, I would know to ask again or check another source.
12	Course navigation support	I would ask the chatbot basic questions about the course (e.g., how to find the syllabus, when the teachers' office hours are)
13	Likely use	I would use this chatbot if I got stuck and a teacher was not available right away.
14	Likely use	I think other students in my class would likely use the chatbot to complete their course assignments.

A Preliminary Semantic Screen for IRT Local Dependence

Qian Shen
College of Education
University of Florida
Gainesville, FL 32611, USA
qian.shen1@ufl.edu

Abstract

This work-in-progress study examines whether sentence embeddings can flag item pairs at risk for local dependence in multidimensional IRT. Using a 50-item personality inventory, preliminary analyses link semantic similarity to residual item-pair dependence. Future work will test robustness across larger samples, additional datasets, and alternative diagnostics.

1 Introduction

As MBTI, Big Five, and other personality-type testing standards have received increasingly widespread attention and recognition, people have begun to try to understand and pay sustained attention to their own corresponding personality characteristics, and they often assess their personality type by completing a wide variety of psychological test items. As with all testing items, researchers also frequently use IRT to analyze personality-type measurement items. Personality-type measurement items typically need to cover multiple dimensions; for example, MBTI test items need to cover 4 dimensions, whereas Big Five personality test items need to cover 5 dimensions. Traditional IRT assumes conditional independence among items, but in practice, this assumption is often violated due to semantic overlap among items, resulting in LD. Within each dimension and across dimensions of personality-type measurement item sets, LD may arise because of similarities in the wording and text of item stems, so that when people answer different items that appear similar, their responses are influenced by the semantic similarity among these items, and this interference can in turn affect the credibility of the resulting personality test outcomes.

2 Background

Researchers have explored more effective methods for LD detection in IRT models from a statistical-methods perspective using multiple approaches. Tang et al. (2020) proposed a threshold autoregressive item response theory (TAR-IRT) model that additionally accounts for sequential dependence among each examinee's item responses. Liu and Maydeu-Olivares (2013) introduced a new general statistic R^2 to evaluate sources of model misfit and compared its performance with other statistics via simulation studies. Edwards et al. (2018) proposed a new LD diagnostic method based on parameter stability under an item-level jackknife resampling procedure.

Wilson (1988) proposed a method for detecting and interpreting deviations from the inter-item LD assumption for items that may share common stimulus materials or other substantive features. Rajlic (2019) found a significant correlation between violations of the local independence assumption defined by LD strength and the magnitude of bias in item parameter estimates. Moreover, some researchers have explored solutions when LD occurs. Ip (2010) argued that multidimensional response data don't necessarily require multidimensional models and discussed situations in which a unidimensional model with LD may be appropriate as an alternative. Olsbjerg and Christensen (2015) proposed a method for modeling LD in longitudinal 2PL models. This method is based on the concept of item splitting and enables more accurate estimation of changes in the latent variable. However, researchers have thus far primarily devoted their efforts to theoretical exploration of statistical methods, without sufficiently focusing on how to integrate the potential textual similarity among items with LD.

However, researchers have thus far primarily devoted their efforts to theoretical exploration of statistical methods, without sufficiently focusing on how to integrate the potential textual similarity among items with LD.

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 9–16
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

3 Method

3.1 Data Sources

The data I used comes from a five-point Likert Big Five personality scale that comprises 50 items and is intended to measure the five dimensions of the Big Five personality traits. Each item employed a five-point scoring format ranging from "strongly disagree" (1) to "strongly agree" (5). After missing values had been removed, 5,000 samples were selected via random sampling without replacement, and these constitute the data used in this study. This dataset serves as an initial methodological test case rather than a basis for universal screening cutoffs.

3.2 Semantic Feature Extraction

The objective of semantic feature extraction is to convert the text of the personality test items in my data into computable numerical representations, so that the semantic similarity between items can be objectively quantified. The pretrained language model Sentence-BERT (SBERT) (Reimers and Gurevych, 2019) can map sentences of arbitrary length into a fixed-dimensional dense vector space, such that semantically similar sentences are closer together in that space. This study used the pretrained `paraphrase-multilingual-MiniLM-L12-v2` model as the SBERT embedding model, extracted vector representations for the 50 items, and computed the cosine similarity between every pair of items.

3.3 Multidimensional IRT Modeling

IRT provides a rigorous measurement framework that uses latent trait models to describe the mathematical relationship between the probability of response of a respondent and their latent ability. One of the core assumptions of IRT is local independence, namely that, conditional on the respondent's latent trait vector $\boldsymbol{\theta}$, item responses should be independent, as stated in Equation (1).

The violation of this assumption is referred to as LD and typically stems from content overlap among items, common method effects, or correlations in measurement error. Before LD detection is performed, it is essential to first establish a well-fitting IRT model because LD identification is based on residual correlations after controlling for the latent trait. If the baseline model itself fits poorly, the residuals will be confounded by model misspecification, preventing accurate identification of true LD. In this study, I first determined the latent dimensional structure of the data using exploratory factor analysis (EFA), and the results of parallel analysis suggested extracting eight factors.

Based on the eight-dimensional structure, I fitted an eight-dimensional Graded Response Model (GRM) (Samejima, 2016). For item i, let $\mathbf{a}_i$ denote its eight-dimensional discrimination vector and b_{ik} the threshold for category boundary k. For respondent j, $\boldsymbol{\theta}_j$ denotes the eight-dimensional latent trait vector. The category probability is defined in Equation (2), with $P^*_{ij,1} = 1$ and $P^*_{ij,K+1} = 0$ defining the category boundaries.

To avoid the computational challenges faced by parameter estimation in multidimensional IRT models, this study adopted the Quasi-Monte Carlo Expectation-Maximization (QMCEM) (Jank, 2005) algorithm, which uses quasi–Monte Carlo methods to approximate integration, reducing computational complexity from exponential to linear. QMCEM substantially increases computational efficiency while maintaining estimation accuracy, supporting estimation of the eight-dimensional model used in this study. The eight-dimensional GRM fitted in this study achieved excellent fit indices (CFI = 0.963, TLI = 0.948, RMSEA = 0.043), markedly superior to the unidimensional model (CFI = 0.380, TLI = 0.354, RMSEA = 0.154); this not only supports the reasonableness of the multidimensional structure but also provides a reliable basis for subsequent LD analysis.

After the model converged, I used oblimin oblique rotation to obtain a simple-structure factor loading matrix. By examining each item's loadings on 8 dimensions, I assigned each item to the dimension with its highest loading; this operational assignment was used only to classify item pairs as within-dimension or between-dimension and does not imply exclusive loadings. The 8 dimensions each contained 2 to 8 items, and the average loading of items on their primary dimension was 0.62 (SD = 0.18). In addition, I used EAP to estimate each respondent's latent trait scores on the 8 dimensions.

3.4 Local Dependence Detection via Partial Correlation

Traditional methods for detecting local dependence within the IRT framework include the LD χ^2 statistic proposed by (Chen and Thissen, 1997) and the (Yen, 1984) Q_3 statistic. LD χ^2 is based on deviations between observed frequencies and model-

expected frequencies, conducting a chi-square test on a two-way joint frequency table for item pairs. However, in multidimensional models, computing expected frequencies requires integration over the high-dimensional latent trait space, and as the number of categories increases, the computation becomes extremely complex. In view of the limitations of traditional methods, this study employed partial correlation to detect local dependence. Partial correlation measures the remaining linear association between two variables after controlling for a set of variables, and its theoretical basis is fully consistent with the conditional independence assumption in IRT.

Expected a posteriori (EAP) scores were estimated on the eight fitted dimensions. For each item i, responses X_i were regressed on these eight EAP scores, producing residuals e_i. For an item pair (i, j), residual dependence was summarized by the residual correlation in Equation (3). This quantity represents the association after controlling for estimated factor scores, rather than conditioning on error-free latent traits.

To place the residual correlations on a common large-sample screening scale, I applied Fisher's transformation and defined the exploratory statistic in Equation (4). This statistic is used as a Fisher-z screening quantity rather than as the standard Yen Q_3 or Chen–Thissen LD statistic. Preliminary meaningful LD required both $|r_{ij\cdot\theta}| > 0.10$ and $T_{ij} > 10$. In the present sample, every pair satisfying $|r_{ij\cdot\theta}| > 0.10$ also satisfied $T_{ij} > 10$, so the effect-size criterion determined the final meaningful-LD count. Each item pair therefore had a continuous residual correlation, an exploratory T_{ij} score, and a binary preliminary meaningful-LD indicator.

3.5 Linking Semantic Similarity to Local Dependence

After completing semantic feature extraction and local dependence detection, I used correlation analysis to establish the basic association between semantic similarity and LD. I computed Pearson and Spearman correlations. To describe heterogeneity by dimension status, I computed Pearson and Spearman correlations separately for within-dimension and between-dimension item pairs. Because the same item appears in multiple pairs, these pair-level associations were treated as descriptive rather than as independence-adjusted inferential tests.

Second, hierarchical regression analysis was used to describe incremental associations and interaction patterns. The dependent variable was the absolute residual correlation $|r_{ij\cdot\theta}|$. I specified three nested models: Model 1 included only semantic similarity; Model 2 added same-dimension status; and Model 3 added their interaction. The three models are given in Equations (5)–(7). I conducted a simple slope analysis to describe the fitted effect of semantic similarity under the two dimension-status conditions: the within-dimension slope is $\beta_1 + \beta_3$, whereas the between-dimension slope is β_1. Changes in R^2 across the nested models were treated descriptively because item pairs share items and therefore do not constitute independent observations.

Finally, a secondary logistic regression used the preliminary meaningful-LD indicator (0/1) as the dependent variable. The model included semantic similarity, same-dimension status, and their interaction, as shown in Equation (8). Model fit evaluation used $-2\log L$, Nagelkerke pseudo-R^2 (Walker and Smith, 2016), and the Hosmer–Lemeshow goodness-of-fit test (Paul et al., 2013).

4 Results

4.1 Local Dependence Patterns

After controlling for EAP estimates of the eight fitted dimensions, the residual correlations across the 1,225 item pairs had a median of -0.002 (IQR $= [-0.023, 0.021]$), indicating that most pairwise residual associations were close to zero. Using the preliminary meaningful-LD criterion, 120 of 1,225 pairs (9.8%) were flagged. As shown in Table 1, 52.8% of within-dimension pairs and 2.4% of between-dimension pairs met this criterion. The mean absolute residual correlation was also larger for within-dimension pairs ($M = 0.091$, $SD = 0.087$) than for between-dimension pairs ($M = 0.023$, $SD = 0.037$). These pair-level differences are reported descriptively because item pairs are not independent.

4.2 Semantic Similarity Predicts Local Dependence

Semantic similarity showed a preliminary positive association with absolute residual dependence. Across all item pairs, the Pearson correlation between semantic similarity and $|r_{ij\cdot\theta}|$ was $r = 0.500$, while the Spearman correlation was $\rho = 0.271$. Semantic similarity alone accounted for 25.0% of the pair-level variance ($R^2 = 0.250$).

As shown in Table 2, the association was stronger within dimensions ($r = 0.651$, $R^2 = 42.4\%$) than between dimensions ($r = 0.209$, $R^2 = 4.4\%$). Thus, semantic similarity was informative within the operational dimension groups but comparatively weak across groups. These pair-level estimates are descriptive because individual items occur in multiple pairs.

4.3 Hierarchical Regression

Table 3 summarizes the descriptive hierarchical regressions. Semantic similarity alone accounted for $R^2 = 0.250$ of the pair-level variance. Adding same-dimension status increased R^2 to 0.486, and adding the interaction increased it to 0.585 ($\Delta R^2 = 0.100$). As shown in Table 4, the fitted interaction coefficient was $B = 0.386$. The corresponding fitted semantic-similarity slope was 0.054 for between-dimension pairs and 0.440 for within-dimension pairs. These estimates indicate a stronger fitted association within dimensions, but the pair-level uncertainty estimates are not adjusted for repeated item membership.

4.4 Logistic Regression

Table 5 summarizes a secondary logistic regression using the preliminary meaningful-LD indicator as the outcome. Semantic similarity and same-dimension status had positive fitted coefficients, whereas the interaction coefficient was not distinguishable from zero in the unadjusted pair-level analysis ($p = 0.401$). The Nagelkerke pseudo-R^2 was 0.456. Because the same item appears in multiple pairs and the meaningful-LD prevalence was 9.8%, this model is treated as descriptive and in-sample; practical screening performance requires out-of-sample validation.

5 Conclusion

This preliminary analysis suggests that semantic similarity contains useful information about residual item dependence, especially for items assigned to the same dimension. It does not establish a general screening rule and does not replace psychometric calibration or content review. At this stage, the method is best viewed as a way to prioritize item pairs for further investigation.

Because this is an initial analysis of a 5,000-respondent sample from one dataset, robustness across datasets and sample sizes has not yet been established. Future work will evaluate additional datasets, sample sizes, text representations, established LD diagnostics, and dependence-adjusted inference.

References

Wen-Hung Chen and David Thissen. 1997. Local dependence indexes for item pairs using item response theory. *Journal of Educational and Behavioral Statistics*, 22(3):265–289.

Michael C Edwards, Carrie R Houts, and Li Cai. 2018. A diagnostic procedure to detect departures from local independence in item response theory models. *Psychological methods*, 23(1):138.

Edward Haksing Ip. 2010. Empirically indistinguishable multidimensional irt and locally dependent unidimensional item response models. *British Journal of Mathematical and Statistical Psychology*, 63(2):395–416.

Wolfgang Jank. 2005. Quasi-monte carlo sampling to improve the efficiency of monte carlo em. *Computational statistics & data analysis*, 48(4):685–701.

Yang Liu and Alberto Maydeu-Olivares. 2013. Local dependence diagnostics in irt modeling of binary data. *Educational and Psychological Measurement*, 73(2):254–274.

Maja Olsbjerg and Karl Bang Christensen. 2015. Modeling local dependence in longitudinal irt models. *Behavior research methods*, 47(4):1413–1424.

Prabasaj Paul, Michael L Pennell, and Stanley Lemeshow. 2013. Standardizing the power of the hosmer–lemeshow goodness of fit test in large data sets. *Statistics in medicine*, 32(1):67–80.

Gordana Rajlic. 2019. *Violations of unidimensionality and local independence in measures intended as unidimensional: assessing levels of violations and the accuracy in unidimensional IRT model estimates*. Ph.D. thesis, University of British Columbia.

Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084*.

Fumiko Samejima. 2016. Graded response models. In *Handbook of item response theory*, pages 95–107. Chapman and Hall/CRC.

Xiaodan Tang, George Karabatsos, and Haiqin Chen. 2020. Detecting local dependence: A threshold-autoregressive item response theory (tar-irt) approach for polytomous items. *Applied Measurement in Education*, 33(4):280–292.

David A Walker and Thomas J Smith. 2016. Nine pseudo r2 indices for binary logistic regression models. *Journal of Modern Applied Statistical Methods*, 15(1):848–854.

Mark Wilson. 1988. Detecting and interpreting local item dependence using a fannily of rasch models. *Applied psychological measurement*, 12(4):353–364.

Wendy M Yen. 1984. Effects of local item dependence on the fit and equating performance of the three-parameter logistic model. *Applied Psychological Measurement*, 8(2):125–145.

A Equations

$$P(\mathbf{X} \mid \boldsymbol{\theta}) = \prod_{i=1}^{I} P(X_i \mid \boldsymbol{\theta}). \tag{1}$$

$$P(X_{ij} = k \mid \boldsymbol{\theta}_j) = P^*_{ij,k} - P^*_{ij,k+1}, \qquad P^*_{ij,k} = \text{logit}^{-1}\left(\mathbf{a}_i^\top \boldsymbol{\theta}_j - b_{ik}\right). \tag{2}$$

$$r_{ij \cdot \theta} = \text{cor}(e_i, e_j). \tag{3}$$

$$z_{ij}^{(F)} = \text{atanh}(r_{ij \cdot \theta}), \qquad T_{ij} = N \left[z_{ij}^{(F)}\right]^2. \tag{4}$$

$$|r_{ij \cdot \theta}| = \beta_0 + \beta_1 \cdot \text{semantic\_sim}_{ij} + \epsilon_{ij}. \tag{5}$$

$$|r_{ij \cdot \boldsymbol{\theta}}| = \beta_0 + \beta_1 \cdot \text{semantic\_sim}_{ij} + \beta_2 \cdot \text{same\_dimension}_{ij} + \epsilon_{ij} \tag{6}$$

$$|r_{ij \cdot \boldsymbol{\theta}}| = \beta_0 + \beta_1 \cdot \text{semantic\_sim}_{ij} + \beta_2 \cdot \text{same\_dimension}_{ij} + \beta_3 \cdot (\text{semantic\_sim}_{ij} \times \text{same\_dimension}_{ij}) + \epsilon_{ij} \tag{7}$$

$$\begin{aligned} \log\left(\frac{P(\text{LD}_{ij} = 1)}{1 - P(\text{LD}_{ij} = 1)}\right) &= \gamma_0 + \gamma_1 \cdot \text{semantic\_sim}_{ij} + \gamma_2 \cdot \text{same\_dimension}_{ij} \\ &\quad + \gamma_3 \cdot (\text{semantic\_sim}_{ij} \times \text{same\_dimension}_{ij}) \end{aligned} \tag{8}$$

B Tables

Table 1: Local Dependence by Dimensional Structure

Criterion	Within-Dimension ($n = 180$)	Between-Dimension ($n = 1{,}045$)	Ratio
Fisher-z screen ($T_{ij} > 10$)	148 (82.2%)	171 (16.4%)	5.0×
Residual effect size ($\lvert r_{ij\cdot\theta}\rvert > 0.10$)	95 (52.8%)	25 (2.4%)	22.0×
Preliminary meaningful LD (both criteria)	95 (52.78%)	25 (2.39%)	22.1×

Table 2: Correlation between Semantic Similarity and Local Dependence by Item Pair Type

Item Pair Type	n	Semantic vs $\lvert r_{partial}\rvert$	Semantic vs T_{ij}	Variance Explained
Within-dimension	180	$r = 0.651$	$r = 0.614$	42.4%
Between-dimension	1,045	$r = 0.209$	$r = 0.225$	4.4%

Note. Values are descriptive pair-level estimates; no independence-adjusted significance tests are reported.

Table 3: Hierarchical Regression Predicting Local Dependence Intensity

Model	Predictors	R^2	ΔR^2
1	Semantic similarity	0.250	—
2	+ Same dimension	0.486	0.236
3	+ Semantic × Dimension	0.585	0.100

Note. $n = 1{,}225$ item pairs. Estimates are descriptive because item pairs share items.

Table 4: Regression Coefficients for Model 3 (Interaction Model)

Predictor	B	SE	β	t	p
Intercept	0.013	0.003	—	4.48	< 0.001
Semantic similarity	0.054	0.011	0.104	5.95	< 0.001
Same dimension	-0.047	0.008	-0.141	-5.73	< 0.001
Semantic × Dimension	0.386	0.022	0.548	17.14	< 0.001

Note. SE, t, and p values are unadjusted for repeated item membership and are reported as exploratory pair-level statistics.

Table 5: Logistic Regression Predicting Meaningful Local Dependence

Predictor	B	SE	z	p
(Intercept)	-6.448	0.663	-9.719	< 0.001
Semantic similarity	9.408	1.850	5.084	< 0.001
Same dimension	3.935	0.861	4.572	< 0.001
Semantic × Dimension	-2.011	2.396	-0.839	0.401

Note. SE, z, and p values are unadjusted for repeated item membership and are reported as exploratory pair-level statistics.

C Data Availability

Link of the data https://www.kaggle.com/datasets/tunguz/big-five-personality-test

Evaluating Multi-Phase and Granular Strategies for Evidence-Anchored Automated Scoring

Alessia Marigo amarigo@wisc.edu
Laura Wright laura.wright@wisc.edu
Linda Malkin lmmalkin@wisc.edu
Xin Xie xin.xie@wisc.edu

School of Education, University of Wisconsin–Madison
1025 W Johnson St, Madison, WI 53706-1706

Abstract

This study evaluates few-shot large language models (LLMs) on middle-school geoscience responses (N=86 responses $\times$ 11 indicators), separating *presence* agreement from *extract* agreement. Joint prompting and annotation-like rubric guidance plus examples yield the clearest gains; lengthy rubric rewriting and clause-level parsing do not reliably improve extract alignment.

1 Introduction

Large language models (LLMs) can score assessments with high internal reliability (Yavuz et al., 2025), yet accuracy sufficient to generate personalized formative feedback remains challenging (Wisniewski et al., 2020). This is especially important in science education, where multidimensional standards target growth across ideas, practices, and crosscutting concepts (NGSS Lead States, 2013). Effective feedback must speak to each dimension separately.

Mixed LLM pipelines and chain-of-thought scoring (Lee et al., 2024), evidence-centered design for AI scoring (Ercikan and McCaffrey, 2022), multi-agent component extraction (Wang et al., 2026), consistency analyses of LLM scorers (Xue et al., 2026), and rubric-verbosity / inference-level effects (Kucia et al., 2026) motivate our experiments. Annotation-aware essay scoring (Ormerod, 2025) further motivates testing finer response units (clauses), which we evaluate and ultimately do not retain as a primary method.

In prior work (Marigo et al., 2026), fine-tuned Llama 3 showed high reliability (α=0.92) with human raters, while Gemini was competitive on presence κ but weak on Fuzzy κ for spans. Both families could often detect indicator presence yet failed to recover the human consensus evidence span (verbatim snippet). For feedback, the model must be accurate in *evidence anchoring*, not only in presence classification. This paper maps which changes—architecture, rubric engineering, or granularity—actually move that needle.

2 Aims

- **System architecture:** Does a dual-phase detect-then-extract pipeline outperform single-phase end-to-end scoring for presence and for verbatim evidence?
- **Rubric engineering:** Which rubric revisions (error polarity, oracle-informed language, additional shots) improve agreement across indicators?
- **Granularity:** Does clause-level parsing improve evidence alignment with the human consensus (oracle)?

3 Sample

Data come from Grade 7–8 science students in the ONPAR project (2019–2020) (Wright et al., 2022). Four teachers piloted an earth-science writing task (two suburban Midwest; two rural East Coast). Two expert raters developed an 11-indicator three-dimensional science rubric and double-coded 100 responses with a *span* or *none* per indicator, then formed a consensus set.

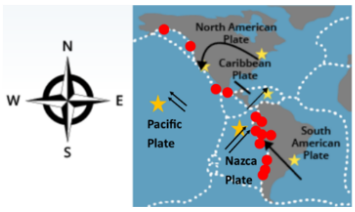
3.1 Writing prompt

The writing prompt provided students with a map of North and South America, tectonic plates in the region, and a series of red dots indicating earthquakes. In addition, a data table provided information about plate movement, including the speed and direction of plate movement. Students were asked to study the information to identify a pattern and make a conclusion about plate movement and earthquakes (see Figure 1).

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 17–23
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

Writing Prompt

The image below shows major earthquakes in a small area of Earth. Study the image and read the data on plate movement. What does the pattern of data allow you to conclude about plate movement and earthquakes?



city	plate	speed mm/ year	north/ south velocity mm/ year	east/ west velocity mm/ year
Port-au-Prince, Haiti	Caribbean	15.45	10.43 N	11.40 E
New York, USA	North American	15.45	4.12 N	14.89 W
San Francisco, USA	North American	18.70	15.22 S	10.87 W
Sao Paulo, Brazil	South American	12.65	12.58 N	1.36 W
Galapagos Islands, Ecuador	Nazca	51.54	8.59 N	50.82 E
Honolulu, USA	Pacific	73.70	24.93 N	69.35 W

Figure 1: ONPAR earth-science constructed-response task: map and table stimuli and the student writing prompt.

3.2 Rubric indicators

The rubric accounts for the three dimensions of the NGSS (NGSS Lead States, 2013): Disciplinary Core Ideas (DCI), Science and Engineering Practices (SEP), and Crosscutting Concepts (CCC). Table 1 lists the 11 indicators used in scoring (canonical names under polarity-inverted rubrics B/C; rubric A uses error-*present* wording for the four error/absence indicators).

3.3 Coding and analysis sample

Of the 100 consensus-coded responses, some were reserved as few-shot examples and some were dropped for near-total *none* coding. The analysis set—hereafter the *oracle*—is **86 responses** $\times$ **11 indicators** (946 cells).

4 Methodology

4.1 Measures

Presence is binary: span $= 1$, none $= 0$. Let $[a, b]$ be the *combined presence codes* for model and oracle.

When both are present ($[1, 1]$ combined presence code), we score span quality. Two spans are a **76% match** if *token-sort* similarity is ≥ 76 (not raw character overlap). We compute this using RapidFuzz (Bachmann, 2025), which splits each span into individual words (tokens) and sorts them alphabetically before calculating their similarity score. This ensures that minor variations in word order or punctuation do not artificially penalize the match. For end-to-end Fuzzy-76 κ we use *combined span codes* (a, b) where:

- $(0, 0)$: both have *none*;
- $(0, 1)$: oracle has a *span* and the model is either *none* or not a 76% match;
- $(1, 0)$: only the model returns a *span*;
- $(1, 1)$: model and oracle both return a *span* with a 76% match.

Prior work yielded usable presence κ but weak Fuzzy-76 κ. We therefore re-ran with the strongest prior configurations—Gemini 3.1-pro (few-shot) and Llama-3.1-8B-Instruct LoRA (r=64, α=16, 5-fold)—and confirmed the gap (Table 2).

4.2 Strategies

4.2.1 Two-phase coding

We separate *detect* (0/1 per indicator) from *extract* (verbatim span only if detect=1). For extract quality we report the agreement rate among eligible cells, that is both oracle and detect return $=1$:

$$76\% \text{ Rate} = \frac{\#(1,1)}{\#(1,1) + \#(0,1)}, \qquad (1)$$

where here $\#(1, 1)$ means 76% match and $\#(0, 1)$ means a mismatch among eligible cells (combined presence code $= [1, 1]$).

4.2.2 Additional measures

Many cells are *none*, partly because several indicators are error/absence constructs: even a strong response yields several zeros, producing imbalance. Cohen's κ is known to be deflated under skewed marginals; indicator-level κ ranged roughly from 0.10 to 0.95. Gwet's AC1 is more stable under imbalance (illustrative range $\approx$ 0.58–0.95 where paired κ was lower).

4.2.3 Rubric revisions

- **A:** original rubric (error indicators scored as error-*present*).
- **B:** A with inverted error polarity (absence-of-error framed as the positive; Table 1).
- **C:** B plus oracle-informed indicator language (what to quote / not quote / when none) and additional paraphrased shots targeting hard extract indicators.

ID	Dim.	Focus (consensus evidence)
R1_INTERACTION	DCI	Two plates interacting; boundaries, borders, or edges
R1_MOVE	DCI	Plate motion (converge/push/collide/slip/transform; relative direction or speed)
R1_HAZARD	DCI	Explicit causal link: earthquakes caused by plate movement
R1_NO_MINOR_ERROR	DCI	Absence of minor motion-wording / out-of-scope error (else quote the error)
R1_NO_MAJOR_ERROR	DCI	Absence of major wrong-motion / wrong-cause error (else quote the error)
R2_DATA_MAP	SEP	Relevant map features (red dots, boundaries, arrows, named plates)
R2_DATA_TABLE	SEP	Relevant table speeds/directions (esp. Nazca / South America)
R2_DATA_INTERPRET	SEP	Correct interpretation or comparison of identified data
R2_NO_INCORRECT_DATA	SEP	Absence of incorrect map/table reading (else quote the error)
R3_PATTERN_IMPLICIT	CCC	Meaningful pattern without an explicit quantifier
R3_PATTERN_EXPLICIT	CCC	Meaningful pattern with explicit quantifier (most/all/preponderance, etc.)

Table 1: Eleven NGSS-aligned scoring indicators. Dimension follows the ONPAR R1/R2/R3 grouping (DCI / SEP / CCC). Brief focus summarizes what consensus raters treat as positive evidence.

Model	Cohen κ	Fuzzy 76% κ
Gemini	0.52	0.25
Llama	0.32	0.07

Table 2: Baseline presence vs. Fuzzy-76 agreement (means over indicators).

4.2.4 Loop by indicator

We also run joint (all indicators in one call) vs. indicator-looped calls for both one-phase and two-phase designs.

4.2.5 Other strategies (tested, not retained)

Since Llama showed frequent hallucinations on this task, we reported arms use Gemini 3.1-pro only. A high-recall detect patch was attempted and while it reduced misses, the number of false alarms exploded. Also a clause-level extract (select or aggregate clauses) was tried but it did not beat thin free-span extract; see Results (Aim 3).

4.2.6 Cont&76% Rate diagnostic

Many mismatches are incomplete but topical: the model span is often a substring (or token subset) of the oracle claim (LLM $\subset$ oracle $\approx$ 50–60% of Fuzzy-76 mismatches on joint arms), while oracle $\subset$ LLM is rarer ($\approx$ 10–15%). We therefore also report

$$\text{Cont\&76\% Rate} = \frac{\#\left\{\begin{array}{l}\text{76\% match}\\ \text{or bidirectional}\\ \text{containment}\end{array}\right\}}{\#\{\text{eligible cells}\}}, \quad (2)$$

where eligible cells are those with combined presence code $= [1, 1]$, as a *diagnostic* of grain (not as the operational scoring rule). Arm-level Cont&76% Rate and 76% Rate are **pooled** over eligible cells.

4.3 Experimental grid

Gemini 3.1-pro; rubrics A/B/C; architectures:

0 one-phase, joint;

1 one-phase, indicator-looped;

2 two-phase detect→extract, joint;

3 two-phase, indicator-looped.

Arm names combine rubric and architecture (e.g., A0, C2).

4.4 Statistical power

With N=86 responses as the clustering unit (946 cells are not independent), minimum detectable differences for paired arm contrasts are on the order of $\approx$ 12–15 percentage points for 76% Rate and $\approx$ 5–12 for presence agreement (80% power, α=0.05, moderate within-response correlation). The size of our dataset is large enough for major metric shifts such as Cont&76% Rate vs. 76% Rate ($\approx$ +30 pp) and some architecture contrasts such as joint rubric vs indicator-looped. It is underpowered for extract improvements based on fine rubric rewording with observed mean deltas of 2-8 pp on 76% Rate.

5 Results

Mean differences across arms in Table 3 look modest, but indicator-level spread is large. Table 4

Arm	Description	Cohen κ	Gwet AC1	76% Rate	Cont&76% Rate
A0	Joint Rubric A 1-Phase	0.528	0.700	0.469	0.823
A1	IndLoop Rubric A 1-Phase	0.466	0.495	0.462	0.817
A2	Joint Rubric A 2-Phases	0.548	0.705	0.441	0.795
A3	IndLoop Rubric A 2-Phases	0.455	0.491	0.492	0.805
B0	Joint Rubric B 1-Phase	0.540	0.698	0.443	0.814
B1	IndLoop Rubric B 1-Phase	0.471	0.524	0.451	0.804
B2	Joint Rubric B 2-Phases	0.547	0.701	0.455	0.809
B3	IndLoop Rubric B 2-Phases	0.472	0.526	0.475	0.792
C0	Joint Rubric C 1-Phase	0.509	0.710	0.485	0.818
C1	IndLoop Rubric C 1-Phase	0.530	0.600	0.502	0.779
C2	Joint Rubric C 2-Phases	0.522	0.720	0.479	0.799
C3	IndLoop Rubric C 2-Phases	0.517	0.578	0.515	0.786

Table 3: Presence and extract metrics by arm. Cohen κ and Gwet AC1 are means over 11 indicators. 76% Rate and Cont&76% Rate are *pooled* over eligible cells (oracle present $\cap$ detect=1).

reports the same pooled Cont&76% Rate as Table 3, with min/max across the 11 indicators (AC1 and Cont&76% Rate).

5.1 Hard indicators and annotation priorities

Presence can look strong while extract fails. Averaging 76% Rate across joint arms A0/B0/C0, the weakest extract indicators are `R1_INTERACTION` ($\approx$ 0.22), `R1_NO_MAJOR_ERROR` ($\approx$ 0.35), `R1_MOVE` ($\approx$ 0.38), `R2_DATA_INTERPRET` ($\approx$ 0.39), and `R3_PATTERN_EXPLICIT` ($\approx$ 0.42). By contrast, `R1_HAZARD` and `R2_DATA_TABLE` often exceed $\approx$ 0.75–0.80 on 76% Rate. `R2_DATA_MAP` is mixed: presence is moderate, but mismatches are frequently bare plate names instead of the full location claim (high Cont&76%, low 76% Rate on several arms).

For future rubric *annotations*, priority should go to:

- **`R1_INTERACTION` / `R1_MOVE`:** require the full interaction or relative-motion claim; forbid quoting only "edges," "boundaries," or bare speeds.
- **`R2_DATA_MAP` vs. `R2_DATA_INTERPRET`:** location/map features vs. comparison/meaning; forbid reusing one span for both.
- **`R3_PATTERN_EXPLICIT`:** quote the pattern-conclusion sentence (with quantifier), not only supporting detail.
- **`R1_NO_MAJOR_ERROR`:** when an error is present, quote only the erroneous causal/motion claim (grain still unstable).

`R1_HAZARD` and `R2_DATA_TABLE` need less annotation effort under current defs.

5.2 Answers to the aims

Aim 1 (architecture). Dual-phase detect→extract does *not* clearly outperform one-phase on 76% Rate once detection errors remain (e.g., A2 0.441 vs. A0 0.469). Joint prompting (arms *0/*2) is more stable for **presence** (AC1 $\approx$ 0.70) than indicator-looped prompting (arms *1/*3), where mean AC1 often falls to $\approx$ 0.50 and some indicators approach zero or negative AC1. Two-phase helps mainly as an analytic separation of detect vs. extract error—not as an automatic accuracy win on mean 76% Rate.

Aim 2 (rubric engineering). Polarity inversion (A→B) clarifies error indicators under AC1 more than it lifts mean 76% Rate. Oracle-informed language plus paraphrased hard-indicator shots (B→C) yields modest extract gains on some looped arms (C1/C3 76% Rate) and competitive joint presence AC1 (C2 0.720), without a large uniform mean presence lift. Annotation-like guidance and examples matter more than lengthy rewrite alone; the hard indicators in Table 5 remain the main extract bottleneck.

Aim 3 (granularity). Clause-level extraction did not improve evidence alignment. Consensus spans are typically full claims (~90–100 characters). Clause-*selection* mismatches remained

Arm	AC1	Min	Max	Cont&76% Rate	Min	Max
A0	0.700	0.489	0.886	0.823	0.462	1.000
A1	0.495	−0.061	0.844	0.817	0.562	1.000
A2	0.705	0.465	0.910	0.795	0.615	0.914
A3	0.491	−0.015	0.826	0.805	0.500	0.970
B0	0.698	0.526	0.909	0.814	0.471	0.971
B1	0.524	−0.061	0.844	0.804	0.562	0.941
B2	0.701	0.349	0.926	0.809	0.500	0.938
B3	0.526	−0.086	0.837	0.792	0.500	0.935
C0	0.710	0.481	0.886	0.818	0.583	1.000
C1	0.600	0.090	0.930	0.779	0.500	0.917
C2	0.720	0.568	0.885	0.799	0.600	0.930
C3	0.578	0.008	0.930	0.786	0.556	0.953

Table 4: Spread across indicators. Gwet AC1 columns are the mean / min / max over indicators. Cont&76% Rate is the *pooled* arm rate (identical to Table 3), with min/max of per-indicator Cont&76% Rate.

Indicator	76% Rate	Cont&76%	AC1
R1_INTERACTION	0.222	0.814	0.877
R1_NO_MAJOR_ERROR	0.350	0.850	0.732
R1_MOVE	0.382	0.764	0.622
R2_DATA_INTERPRET	0.393	0.505	0.735
R3_PATTERN_EXPLICIT	0.421	0.724	0.533
R2_DATA_MAP	0.459	0.909	0.622
R3_PATTERN_IMPLICIT	0.522	0.767	0.544
R2_NO_INCORRECT_DATA	0.548	0.865	0.583
R1_NO_MINOR_ERROR	0.710	0.952	0.706
R1_HAZARD	0.757	0.941	0.886
R2_DATA_TABLE	0.802	0.844	0.887

Table 5: Mean 76% Rate / Cont&76% Rate / presence AC1 by indicator on joint arms A0, B0, and C0 (unweighted mean of the three arms). Lower 76% Rate = higher priority for claim-level rubric annotations and contrastive shots.

mostly **LLM** $\subset$ **oracle** (under-quoting), matching free-span grain errors. Clause-*aggregation* flipped errors toward **oracle** $\subset$ **LLM** (over-merging) and lowered 76% Rate. Clause boundaries are not the same unit humans quote; Cont&76% Rate rises because many Fuzzy-76 "misses" are contained fragments of the oracle. The clause-level extraction helped us diagnose the grain gap and the Cont&76% Rate was key to not underestimate working extraction methods.

6 Discussion

Our results offer key insights for rubric design and model instruction. Recent work shows model-written rubrics underperform human standards in rubric-guided evaluation (Zhou et al., 2026); our results similarly favor expert-calibrated indicator language over extensive automated rewrite. Simply expanding a rubric or flipping its polarity does not guarantee automatic gains in extract accuracy. Instead, targeted, expert-informed instructions and precise examples (such as those in Rubric C) are proven to be more effective in addressing specific extraction bottlenecks.

For evidence-anchored formative scoring, we recommend reporting presence and extraction agreement separately, favoring AC1 (or similar) alongside κ under imbalance, and treating Cont&76% Rate as a sensitivity check on grain. Operational systems should prioritize joint presence prompts, claim-level instructions/shots focused on the hard indicators above, and human-like calibration examples—not clause chopping or detect-recall patches that trade misses for false alarms. Given N=86 responses, small mean arm deltas should be interpreted cautiously (see power note above).

At a higher level, our methodology represents a highly promising first step toward scaling automated writing evaluations that are explainable, cost-effective, and accurate. Practically, by pairing binary presence scores with verbatim evidence spans, our approach creates an instant audit trail for automated grading that allows classroom teachers to verify AI-generated feedback at a glance, dramatically lowering the barrier to trust and deployment in live classrooms. Furthermore, our finding that joint prompting outperforms looped architectures translates to an immediate operational benefit, as developers can obtain multidimensional feedback in a single API call rather than N individual queries, reducing processing costs and network latency, and

making real-time, student-facing formative scoring economically and technically viable at scale. Specifically, to optimize both detection stability and extraction precision, we suggest a hybrid, two-step pipeline for operational designs: first, run a one-phase prompt for detection, ignoring the initial spans generated; second, run the second phase of the two-phase prompt to extract evidence for the indicators detected in the first step. This avoids wasting token overhead or model attention on absent concepts, focusing the computational budget entirely on retrieving high-quality, verbatim evidence.

A key theoretical implication of our work is that LLMs cannot simply bypass the deep pedagogical preparation traditionally required for human raters. In high-stakes educational testing, human scorers undergo rigorous training and rely on pages of scoring materials, anchor charts, and concrete examples illustrating when to apply specific codes. We observed that LLMs face a parallel bottleneck: while they easily classify the presence of a rubric concept, they fail to anchor the exact verbatim evidence without this same expert guidance and exemplary context. Consequently, the path to robust automated scoring is not through structural pipeline engineering, but through translating human calibration materials into model instructions. A crucial next step for our research will consist of systematically developing these detailed scoring annotations and empirically comparing LLM performance with and without these comprehensive calibration materials.

Limitations

This study has several limitations. First, the analysis sample is modest (N=86 responses; 946 indicator cells that are clustered within responses), so small mean differences between arms are often underpowered (Section 4.6); larger multi-task samples would be needed to confirm architecture effects. Second, we evaluate a single ONPAR earth-science prompt and one final model family (Gemini 3.1-pro few-shot); results may not transfer to other domains, grade bands, or open-weight models (Llama was dropped after hallucination issues on this task). Third, the oracle is a two-rater consensus with spans that themselves encode a particular evidentiary grain; disagreement with the model is not always "model error" versus an equally defensible alternate span. Fourth, Fuzzy-76 and Cont&76% Rate are operational metric choices: lowering the similarity cutoff or accepting containment raises apparent extract agreement without changing the underlying quotes. Fifth, rubric C annotations and paraphrased shots were informed by error analysis on related runs; while shots paraphrase student language, residual optimism about annotation benefits is possible and should be tested on a fully held-out task. Sixth, we do not evaluate downstream feedback quality with teachers or students—evidence agreement is necessary but not sufficient for formative usefulness. Seventh, we do not report human–human span agreement on the same Fuzzy-76 / Cont&76% metrics, which would provide a ceiling for interpretability. Finally, these samples were collected in 2019 and there have been significant changes in student classroom writing practices since that time. The use of generative AI, in particular, may affect how well the findings apply to current formative assessment.

References

Max Bachmann. 2025. rapidfuzz/rapidfuzz: Release 3.13.0.

Kadriye Ercikan and Daniel F. McCaffrey. 2022. Optimizing implementation of artificial-intelligence-based automated scoring: An evidence-centered design approach for designing assessments for AI-based scoring. *Journal of Educational Measurement*, 59(3):272–287.

Filip J. Kucia, Anirban Chakraborty, and Anna Wróblewska. 2026. LLM essay scoring under holistic and analytic rubrics: Prompt effects and bias. *arXiv preprint arXiv:2604.00259*.

Gyeong-Geon Lee, Ehsan Latif, Xuansheng Wu, Ning Liu, and Xiaoming Zhai. 2024. Applying large language models and chain-of-thought for automatic scoring. *Computers and Education: Artificial Intelligence*, 6:100213.

Alessia Marigo, Laura Wright, and Linda Malkin. 2026. Scoring student science responses with small training datasets: Comparing LLMs and fine-tuned models. Poster presented at the annual meeting of the National Council on Measurement in Education (NCME). Poster presentation.

NGSS Lead States. 2013. *Next Generation Science Standards: For States, By States*. The National Academies Press, Washington, DC.

Christopher Ormerod. 2025. Automated essay scoring incorporating annotations from automated feedback systems. In *Proceedings of the Artificial Intelligence*

in Measurement and Education Conference (AIME-Con): Coordinated Session Papers, pages 9–18. National Council on Measurement in Education.

Yucheng Wang, Zhiyao Ding, Xuansheng Wu, Shengwei Sun, Ning Liu, and Xiaoming Zhai. 2026. AutoSCORE: Enhancing automated scoring with multi-agent large language models via structured component recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pages 40898–40906.

Benedikt Wisniewski, Klaus Zierer, and John Hattie. 2020. The power of feedback revisited: A meta-analysis of educational feedback research. *Frontiers in Psychology*, 10:3087.

Laura Wright, H. Harkins, Rebecca Kopriva, W. Auty, Linda Malkin, and B. Myers. 2022. Technology-enhanced tasks to assess three-dimensional science sense-making: Possibilities and lessons learned from the ONPAR NGSS-based classroom assessment project. *Contemporary Issues in Technology and Teacher Education*, 22(2).

Mingfeng Xue, X. Xiao, Y. Liu, and Mark Wilson. 2026. On the consistency of automatic scoring with large language models. *Educational and Psychological Measurement*. Article 00131644261418138.

Fatih Yavuz, Özkan Çelik, and Gülşah Yavaş Çelik. 2025. Utilizing large language models for EFL essay grading: An examination of reliability and validity in rubric-based assessments. *British Journal of Educational Technology*, 56(1):150–166.

Junyi Zhou, Qiyuan Zhang, Yufei Wang, Fuyuan Lyu, Yidong Ming, Can Xu, Qingfeng Sun, Kai Zheng, Peng Kang, Xue Liu, and Chen Ma. 2026. RubricBench: Aligning model-generated rubrics with human standards. In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 31179–31200, San Diego, California, United States. Association for Computational Linguistics.

Consensus without Accuracy: Investigating LLM's Ability to Recover Item Difficulty Using Paired Comparisons

Michael L. Chrzan
Center for Educational Data Science and Innovation,
University of Maryland

Benjamin W. Domingue
Stanford University *

Abstract

Large language models (LLMs) offer a scalable way to estimate item difficulty before respondent data are available, but it is unclear whether they can provide reliable estimates, either alone or in aggregate. We evaluate seven frontier LLMs on pairwise comparisons of items from seven datasets with respondent-calibrated one-parameter item response theory (1PL IRT) difficulty estimates. We aggregate their judgments using a hierarchical Bayesian Bradley–Terry model that estimates shared item strengths while accounting for model-specific discrimination and presentation-order bias. The pooled rankings recover a meaningful but incomplete difficulty signal: standardized hierarchical Bradley–Terry strengths correlate moderately with IRT difficulty overall (Spearman's $\rho = 0.400$, $p < .001$), with dataset-level correlations ranging from 0.250 to 0.591. This heterogeneity persists despite high inter-model agreement (Krippendorff's $\alpha = 0.834$–0.917) and generally high within-model triplet transitivity, demonstrating that consensus and internal consistency do not guarantee external psychometric validity. The hierarchical model also reveals substantial variation in rater discrimination and position bias, while in-sample predictive accuracy ranges from 0.732 to 0.874 across datasets. These findings support LLM pairwise comparisons as a scalable source of provisional item ordering, but not as a replacement for respondent-based calibration.

1 Introduction

Item calibration - the process of estimating psychometric properties such as item difficulty - is central to educational assessment but resource-intensive, typically requiring large samples of student response data and extended field-testing periods creating barriers for low-resource testing programs, early-stage item development, and rapid test assembly contexts (Jiang et al., 2016). Recent advances in large language model (LLM) capabilities raise an intriguing question: can item difficulty be meaningfully estimated directly from item text, without recourse to any response data? Such an approach would offer immediate practical utility for item writers and test developers seeking a scalable first-pass difficulty signal for item screening, ordering, or rough-cut form assembly before formal calibration. Prior work has explored LLMs for automated item generation and scoring (von Davier et al., 2022), but their deployment as a calibration proxy for estimating the psychometric difficulty of items never yet administered remains less examined as does the extent to which paired comparisons aid in that effort.

Recent work by Wu et al. (2024) introduced concept-guided chain-of-thought (CGCoT) prompting as a generalizable framework for scoring texts on latent constructs. CGCoT pairs a sequence of researcher-crafted prompts, each probing a different facet of the latent construct of interest, with LLM-based pairwise comparisons aggregated into a continuous score via the Bradley-Terry (BT) model (Bradley and Terry, 1952). Originally validated for scaling political affect in social media, CGCoT's structure maps naturally onto item difficulty estimation: an item's text can be characterized through prompts targeting cognitive demand and required ability level, with pairs of items compared by an LLM to establish an ordinal difficulty ranking without any examinee data.

Further, pairwise comparison methods offer a natural interface for eliciting difficulty judgments from LLMs. Rather than requiring absolute estimates - which are psychometrically undefined without a response scale and shown to be unreliable for LLMs (Licht et al., 2025) - pairwise comparisons ask a model to make relative judgments, which can then be aggregated into difficulty rankings. This approach also allows reliability to be assessed via ex-

*Correspondence: mlchrzan@umd.edu

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 24–36
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

tensions of Bradley-Terry methods, such as a Hierarchical Bradley-Terry Model (Böckenholt, 2001), and internal consistency metrics such as triplet transitivity. However, whether LLM-derived BT scores actually align with psychometrically validated difficulty estimates from Item Response Theory (IRT) remains an open empirical question.

This study makes three contributions to answering these questions. First, we provide a multi-model, multi-dataset evaluation of the extent to which LLM pairwise judgments recover respondent-calibrated item difficulty. Second, we aggregate judgments with a hierarchical Bayesian Bradley–Terry model that estimates a shared item-difficulty ranking while separating model-specific discrimination from presentation-order bias, which quantifies uncertainty in the pooled ranking and reveals consequential differences among LLM raters that a simple consensus score would obscure. Lastly, we distinguish three often-conflated properties of LLM judgments: within-model transitivity, inter-model agreement, and external psychometric validity. Despite generally high transitivity and uniformly high inter-model agreement, alignment with 1PL difficulty remains moderate and heterogeneous, showing that consistent LLM consensus is useful for provisional item ordering but cannot substitute for respondent-based calibration.

2 Related Work

The foundational IRT model utilized in large-scale calibration is the 1-Parameter Logistic (1PL) model, commonly known as the Rasch model, introduced by Georg Rasch in 1960 (Rasch, Georg, 1960; Kolesnikova et al., 2026). The Rasch model conceptualizes the probability that a specific student correctly answers a specific item as a logistic function of the difference between the student's latent ability and the item's latent difficulty. In generalized terms, if a student's ability exactly matches the item's difficulty, the theoretical probability of a correct response is precisely 50%. In traditional workflows, estimating this difficulty parameter requires administering the item to highly calibrated populations whose abilities form a known distribution (Felice and Karaoz, 2024). However, when new assessment items are generated, their parameters are entirely unknown, creating what is known as the cold-start problem (Wainer et al., 2000). If LLMs or machine learning ensembles can accurately predict the difficulty parameter based solely on item content, these synthetic estimates can be used to jump-start the IRT calibration process. This permits the immediate deployment of new items in live environments and vastly reduces the volume of human response data required to achieve statistical stability. While prior work has examinined the ability of other families of machine learning models to predict item difficulty (Felice and Karaoz, 2024; Kapoor et al.), only Kolesnikova et al. (2026) used pairwise comparisons as a methodology and their conclusions were limited by the singular bank of items.

An intuitive methodology for applying LLMs to predict item difficulty is direct absolute estimation. In this paradigm, a language model is provided with the text of an item—including the question stem, correct answer, and distractors—and is instructed via a zero-shot prompt to output a numeric difficulty score or a categorical difficulty level (Kolesnikova et al., 2026). This approach relies on the implicit assumption that the LLM contains a holistic, internally calibrated representation of cognitive complexity that perfectly mirrors human populations. Recent literature indicates that this assumption is fundamentally flawed and suggests that direct absolute judgment by LLMs is fraught with inconsistencies and limitations (Licht et al., 2025; Li et al.; Wang et al., 2026).

To overcome the inherent limitations of direct absolute prediction with LLMs, researchers have achieved state-of-the-art results by altering the LLM's role within the predictive workflow. Rather than acting as a final psychometric judge, the LLM is repositioned as a highly advanced feature extractor or simulator (Hoyl; Kolesnikova et al., 2026; Scarlatos et al.). In this paradigm, the model is prompted to evaluate specific linguistic, structural, and pedagogical characteristics of the item text. These extracted qualitative features are then fed into traditional supervised machine learning algorithms to predict the final difficulty parameter. By decoupling the semantic extraction (handled by the LLM) from the regression task (handled by ensemble trees like Random Forests and Gradient Boosting algorithms), the predictive system successfully sidesteps the LLM's inability to globally scale difficulty (Felice and Karaoz, 2024).

While feature-based ensembles and synthetic simulations provide robust solutions, extracting dozens of qualitative features per item can be computationally expensive, highly prompt-intensive, and strictly domain-dependent. An alternative,

highly effective psychometric strategy leverages the language model's fundamental capacity for relative reasoning: pairwise (paired) comparisons. While assigning an isolated assessment item a difficulty score of 7.8 out of 10 requires a stable internal standard, precise calibration against prior examples, and a consistent, unshifting interpretation of what the numerical scale is meant to represent (DiGiuseppe and Flynn, 2026), pairwise comparisons ask a significantly simpler question: Which of these two items is more difficult? This localized, relative judgment completely removes the need for a global numerical frame. It severely mitigates anchoring bias, allows for high-resolution discrimination between items of similar complexity, and demonstrates remarkable stability across various LLM architectures and model sizes (DiGiuseppe and Flynn, 2026). By transforming the task from an absolute regression to a binary relative choice, the system extracts the model's highest-confidence judgments and has been shown to be highly effective in recent work examining mathematics tasks (Kolesnikova et al., 2026).

However, when utilizing different LLMs or deploying varying prompts to generate large-scale pairwise comparisons the reliability of the artificial "judges" inherently fluctuates. Inspired by the established Crowd-BT algorithm (Chen et al., 2013)—which automatically estimates the underlying reliability of human crowdsource workers and flags annotators with low quality for removal—researchers have introduced the BT-σ model, which extends the traditional Bradley-Terry framework specifically for the "LLM-as-a-jury" setting (Qian et al., 2026). The BT-σ Bayesian model introduces a specific discrimination parameter (σ_k) for each individual LLM judge. This architectural enhancement allows the model to jointly infer the global item difficulty rankings while simultaneously weighting the statistical reliability of each judge using pairwise comparisons alone. If a specific language model frequently outputs comparisons that contradict the consensus of the optimized matrix, its parameter is mathematically penalized, permanently down-weighting its future influence on the latent difficulty scores. This creates a highly resilient, noise-resistant ecosystem capable of safely fusing predictions from advanced, proprietary models with smaller, open-weight models, resulting in a single, cohesive, and highly accurate difficulty continuum. Building on this work, we apply rater-aware Bayesian pairwise aggregation to educational item calibration, jointly modeling LLM-specific discrimination and presentation-order/position bias and validating the resulting pooled rankings against respondent-calibrated 1PL IRT difficulty across seven datasets.

3 Data

The items we examine were drawn from seven datasets in the Item Response Warehouse (IRW) (Domingue et al., 2025). Using their IRW tags, they are: (**1**) `frac20` - a mathematics/fractions dataset (Tatsuoka, 2002) - and six literacy-focused datasets from three studies: (**2**) `gilbert_meta_2` (Kim et al., 2023), (**3**) `gilbert_meta_7` and (**4**) `gilbert_meta_8` (Kim et al., 2024), as well as (**5**) `gilbert_meta_102`, (**6**) `gilbert_meta_103`, and (**7**) `gilbert_meta_104` (Relyea et al., 2025). These datasets include items with established 1PL IRT difficulty calibrations, providing a psychometric ground truth for validation. They were selected using the IRW's filtering capabilities to ensure that the datasets have item text available, the age range is <=18, the construct is cognitive/educational, the measurement tool is test, the item format is Likert/constructed response, and language is English.

Seven proprietary LLM configurations served as raters: (**1**) GPT-5.4 Mini, (**2**) GPT-5.4, (**3**) GPT-5.6-Luna, (**4**) GPT-5.6-Sol, (**5**) Gemini 3 Flash Preview, (**6**) Gemini 3.7 Flash, and (**7**) Gemini 3.1 Pro Preview. The panel was selected to vary along three dimensions relevant to the robustness of LLM-based measurement. First, it includes two independently developed model families, OpenAI and Google, reducing the extent to which the findings depend on the idiosyncrasies of a single provider. Second, it includes different tiers and variants within each family, allowing us to examine whether pairwise judgments remain consistent across models marketed for different capability, latency, and cost tradeoffs. Third, the numbered releases within both families broaden the evaluation beyond a single model generation or temporal snapshot. Because the providers do not disclose directly comparable parameter counts, training corpora, or complete architectural details for these proprietary systems, we treat each configuration as a distinct rater rather than interpreting nominal tier or version differences as quantitative measures of model size or capability.

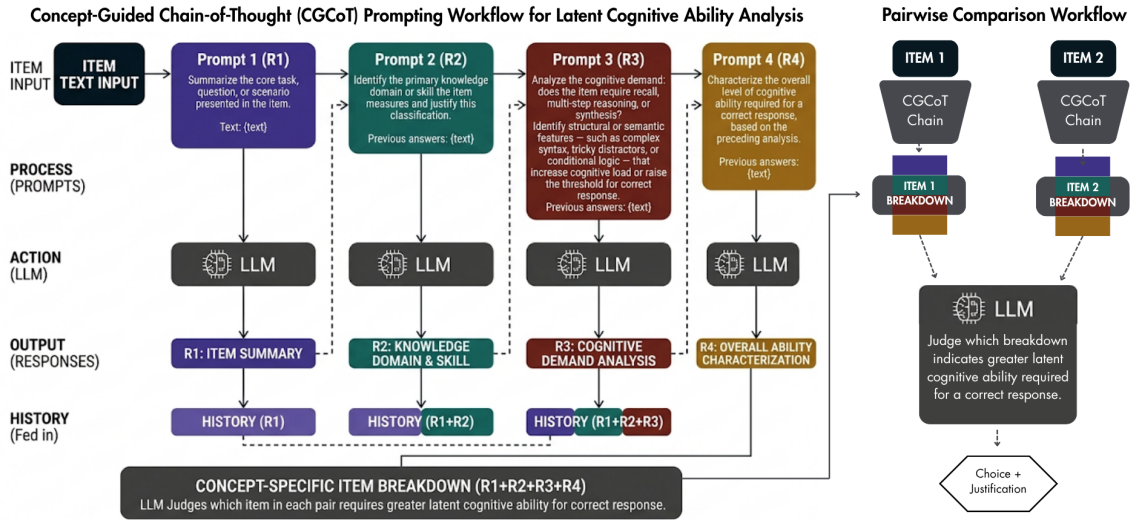


Figure 1: Concept-Guided Chain-of-Thought (CGCoT) Prompting for Item Difficulty Assessment. Each item is processed through a sequence of four prompts, each targeting a different facet of the latent construct of interest, here item difficulty. The LLM's responses are concatenated and fed into the next prompt, culminating in a comprehensive analysis of the item's difficulty before pairwise comparisons are made.

4 Methods

4.1 `pairadigm` and Paired Comparisons

All pairwise comparison workflows—CGCoT prompt chaining, comparison management, and BT model estimation—were implemented using `pairadigm` (Chrzan, Michael L.), an open-source Python package for using LLMs to measure scalar constructs. The package includes code for generating pairs from the set of items and applying CGCoT prompts to construct item breakdowns. Using evidence from Wu et al. (2024) that beyond 10 pairs per item yielded minimal improvement, where possible we limit the number of possible pairs to 10 per item for compute and cost considerations. The seven configurations listed previously were evaluated as separate comparison judges, enabling rater-specific estimates of consistency, discrimination, and position bias.

For each item, a four-prompt CGCoT chain was administered sequentially to the LLM to target cognitive item difficulty. After the LLM answers each prompt, its response is concatenated to all its previous responses and fed into the next prompt until at the end an entire breakdown of the item is created, as displayed in Figure 1. Prompts are also listed out in Appendix A.

The resulting concept-specific breakdown for each item was used to generate pairwise comparisons across all item pairs within each dataset. Based on the cumulative output from these four steps applied to an item, the LLM judged which item in a pair required greater latent cognitive ability for correct response.

4.2 (Hierarchical) Bradley-Terry Models

Individual BT models were fit to each LLM's pairwise preference data for each dataset to produce a continuous BT difficulty score for each item. Triplet transitivity, or the proportion of triplets (A, B, C) in which the model's pairwise judgments were logically consistent, was computed as an additional consistency check. Inter-rater agreement across LLMs was estimated using Krippendorff's Alpha (Hayes and Krippendorff) to assess the consistency across models for a given IRW dataset. Finally, convergent validity was assessed via Spearman rank correlations computed between each model's BT scores and the corresponding 1PL difficulty parameters, stratified by model and dataset (Spearman). Relationships among reliability, inter-rater agreement, and Spearman correlations were examined to assess the relationships between internal consistency and external validity.

A Hierarchical Bradley-Terry model was also fit to assess each LLM's discrimination and position bias. We formulate the model as follows:

$$\text{logit}P(y_{ijk} = 1) = \alpha_k(\lambda_i - \lambda_j) + \gamma_k$$

where

- $y_{ijk} = 1$ indicates that model k selects the first presented item, i, over the second-presented item, j,

- λ_i and λ_j are the difficulty parameters for items i and j,
- α_k is the discrimination parameter for participant k,
- γ_k is the position bias for participant k.

We structure the data such that every unordered pair is evaluated in both orientations, which balances item identity across first and second positions and makes γ_k estimable from direct orientation reversals. Further, we enforce that (1) the sum of λ_i across all items is zero, and (2) the geometric mean of α_k across all models is one, which identifies the global $\alpha - \lambda$ scale without choosing an arbitrary reference model. We discuss the full model details in the Appendix, including the hierarchical priors for λ_i, α_k, and γ_k in Appendix B as well as model diagnostic checks in Appendix C and sensitivity analysis in Appendix D.

5 Results

5.1 Recovery of Respondent-Calibrated Difficulty

Across the 179 items, standardized HBT item strengths were moderately associated with respondent-calibrated 1PL difficulty (Spearman's $\rho = 0.400$, $p < .001$). As Figure 2 and Table 1 show, however, this relationship varied substantially by dataset. The strongest association occurred for `frac20` ($\rho = 0.591$, $p = .006$). Positive associations were also observed for `gilbert_meta_102` ($\rho = 0.449$, $p = .006$), `gilbert_meta_8` ($\rho = 0.423$, $p = .022$), `gilbert_meta_103` ($\rho = 0.414$, $p = .025$), and `gilbert_meta_7` ($\rho = 0.398$, $p = .016$). Estimates were weaker and not statistically distinguishable from zero for `gilbert_meta_2` ($\rho = 0.337$, $p = .146$) and the nine-item `gilbert_meta_104` dataset ($\rho = 0.250$, $p = .516$). Thus, the pooled ranking recovered a meaningful difficulty signal, but not uniformly across item sets. In particular, the strongest result arose for the mathematics dataset, whereas results varied among the six literacy datasets; the observed heterogeneity therefore does not reduce to a simple mathematics-versus-literacy distinction. The corresponding model-specific correlations, pooled across all items and disaggregated by dataset, are reported in Appendix Tables 2 and 3, respectively.

The HBT ranking also retained substantial information from the model-specific BT rankings. Across all items, Spearman correlations between each individual model's BT scores and HBT scores ranged from 0.545 for GPT-5.4 Mini to 0.864 for GPT-5.6 Sol. Hierarchical pooling therefore captured a broadly shared ordering without making any single LLM the reference rater.

5.2 Consistency and Inter-Model Agreement

Most model–dataset combinations exhibited high triplet transitivity, often exceeding 95% and matching transitivity observed in human judgments (Sleboda and Sokolowska). The principal exceptions involved GPT-5.4 Mini, whose transitivity was 76.2% on `gilbert_meta_104`, 87.9% on `gilbert_meta_2`, 91.4% on `gilbert_meta_7`, and 91.9% on `gilbert_meta_102`; GPT-5.4 also reached only 85.7% on `gilbert_meta_104`. Complete model-by-dataset estimates and triplet counts appear in Appendix Table 4. Agreement across the seven LLMs was uniformly high by conventional standards, with dataset-level Krippendorff's α ranging from 0.834 to 0.917. These strong internal-consistency and inter-model-agreement results contrast with the moderate and heterogeneous correlations with 1PL difficulty. Transitivity, agreement among LLMs, and alignment with respondent-calibrated difficulty consequently capture distinct properties of the judgments.

5.3 Rater Effects and Model Fit

Figure 3 shows pronounced variation in both rater discrimination and presentation-order bias. Posterior median discrimination estimates ranged from 0.210 for GPT-5.4 Mini on `gilbert_meta_7` (94% HDI $[0.137, 0.285]$) to 6.013 for Gemini 3.7 Flash on `gilbert_meta_104` (94% HDI $[2.399, 20.124]$); the wide interval for the latter estimate reflects the small number of items. GPT-5.4 Mini generally had low discrimination on the Gilbert datasets. Position effects were likewise consequential. The most negative estimate was for GPT-5.4 Mini on `gilbert_meta_104` ($\gamma = -1.927$, 94% HDI $[-2.678, -1.254]$), indicating a preference for the second-presented item. The largest positive estimate occurred for Gemini 3 Flash Preview on `gilbert_meta_8` ($\gamma = 0.858$, 94% HDI $[0.599, 1.126]$), indicating a preference for the first-presented item. Because each unordered pair appeared in both orientations, these tendencies could be separated from item identity and shared item strength.

The HBT item ordering agreed with observed

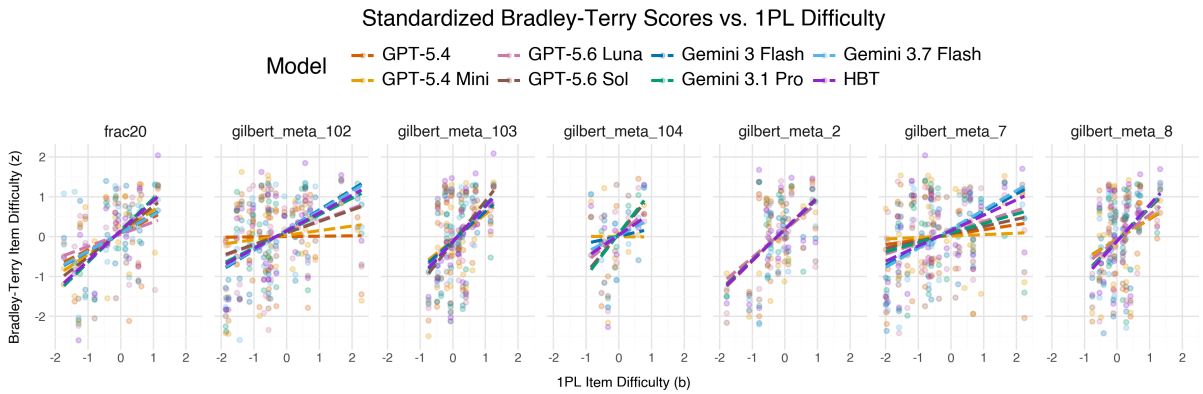


Figure 2: Standardized Bradley–Terry scores versus 1PL IRT difficulty. Each panel represents one dataset; points and dashed linear trends distinguish the seven model-specific BT estimates and the pooled HBT estimate. The pooled HBT ranking is positively associated with respondent-calibrated difficulty overall, but the strength of association varies across datasets and models.

Dataset	Items	Judgments	ρ	p	Krip. α	Rank agr.	Accuracy	Brier
`frac20`	20	2,324	.591	.006	.839	.732	.732	.174
`gilbert_meta_2`	20	2,310	.337	.146	.899	.834	.853	.097
`gilbert_meta_7`	36	4,200	.398	.016	.864	.784	.794	.136
`gilbert_meta_8`	29	3,430	.423	.022	.902	.848	.854	.101
`gilbert_meta_102`	36	4,340	.449	.006	.850	.760	.778	.148
`gilbert_meta_103`	29	3,430	.414	.025	.917	.871	.874	.087
`gilbert_meta_104`	9	546	.250	.516	.834	.716	.787	.136

Table 1: Dataset-level validity, agreement, and HBT fit. ρ is the Spearman correlation between standardized HBT item strength and 1PL difficulty; α is Krippendorff's inter-model agreement coefficient. Rank agreement is the posterior mean agreement between the deterministic HBT item ordering and observed winners. Accuracy and Brier scores use the full HBT likelihood and are in-sample. Judgments count model-specific comparisons across both pair orientations.

pairwise winners in 71.6% to 87.1% of judgments across datasets. When rater-specific discrimination and position bias were included through the full likelihood, in-sample predictive accuracy ranged from 73.2% on `frac20` to 87.4% on `gilbert_meta_103`. The clearest gain from incorporating rater effects occurred for `gilbert_meta_104`, where agreement based on item ordering alone was 71.6%, compared with 78.7% full-likelihood predictive accuracy. These values describe fit to the judgments used for estimation and are not measures of performance on unseen comparisons. Full sampling and predictive diagnostics are reported in Appendix C.

6 Discussion

The central result is qualified but substantively useful: LLM pairwise comparisons contain recoverable information about psychometric item difficulty, yet that information is incomplete. Pooling seven LLMs in the HBT model produced an overall correlation of 0.400 with 1PL difficulty, demonstrating that judgments based on item text alone capture more than arbitrary preference. At the same time, the dataset-level range from 0.250 to 0.591 is too broad, and the weaker results too consequential, to treat LLM-derived rankings as interchangeable with respondent-based calibration.

The contrast between consensus and validity is particularly important. The LLMs usually produced transitive judgments and agreed with one another, but their shared ranking aligned only moderately with IRT difficulty. Agreement among artificial raters can therefore indicate a stable model-derived construct without establishing that this construct is the one measured in human response data. Models may converge on salient textual cues, curricular content, or solution complexity while missing factors that shape empirical difficulty. Increasing the number of similar raters may sharpen such a shared ordering without correcting a common construct mismatch.

Additionally, the rater parameters we estimate help explain why hierarchical aggregation is prefer-

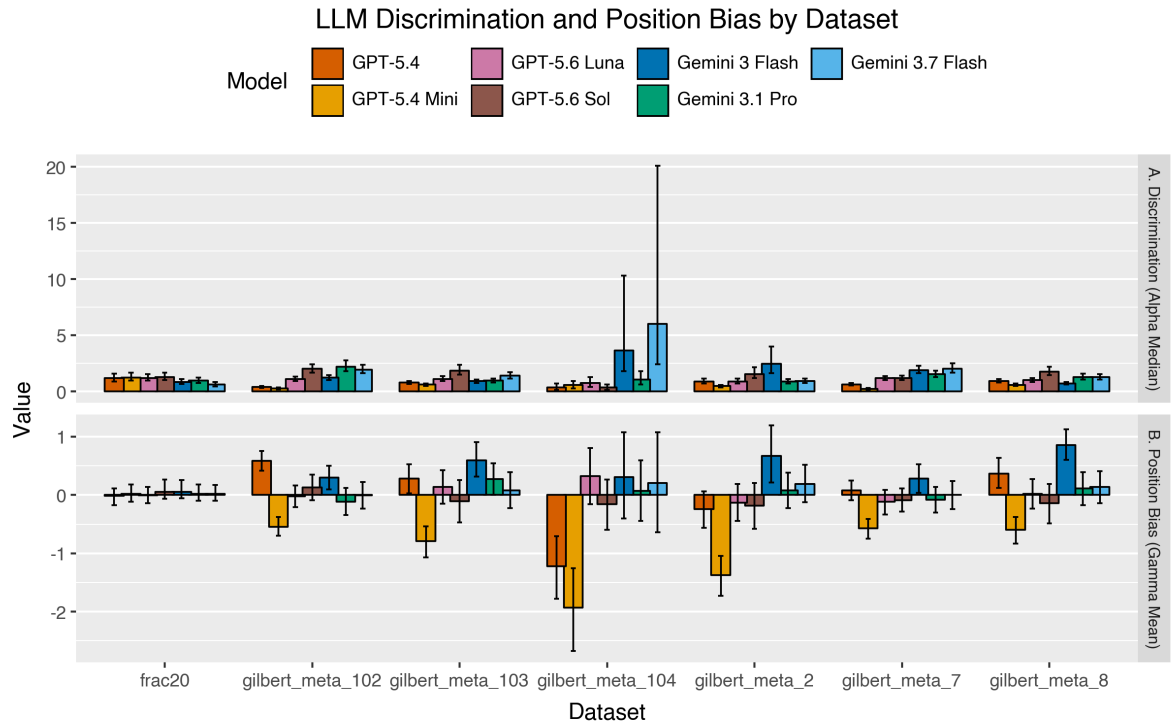


Figure 3: Discrimination and Position Bias of LLMs. The figure shows the discrimination (Alpha) and position bias (Gamma) of different LLM raters across various datasets. Higher alpha values indicate stronger rater ability to distinguish strength between items, while positive gamma values suggest a bias toward the first item shown. Error bars represent 94% highest density intervals (HDI) for each estimate. The results highlight the variability in LLM performance across datasets, with some models exhibiting higher discrimination and lower position bias, suggesting better reliability and less susceptibility to order effects in pairwise comparisons.

able to selecting one model or averaging rankings without adjustment. LLMs differed markedly in how sharply they responded to latent item-strength differences, and several showed presentation-order effects large enough for their 94% intervals to exclude zero. More generally, the results argue for treating LLMs as heterogeneous judges whose decisions require measurement, rather than as exchangeable replications of a single evaluator.

Performance differences across datasets also caution against broad domain-level claims. The one mathematics set in the IRW with item text available at time of writing showed the strongest external alignment, but the literacy datasets ranged from weak to moderate despite their related content. Item count contributes to uncertainty, especially for `gilbert_meta_104`, but sample size alone cannot explain the heterogeneity. Differences in item format, construct representation, response processes, and the correspondence between textual demand and empirical difficulty are plausible sources that should be tested directly as access to item text grows.

For assessment development, these findings support a limited but practical role for LLM-derived rankings. Considering our results using paired comparisons with those use direct absolute prediction, the evidence suggests LLMs can provide immediate provisional ordering for item review, rough form assembly, or prioritization before respondent data are available but they should not be used as final calibration parameters or as a substitute for field testing. Stronger applications would combine balanced presentation, multiple model families, inspection of rater-specific α_k and γ_k, and anchor items with known IRT parameters. Validation should ultimately be conducted on held-out comparisons and new item sets, with respondent calibration retained as the criterion.

Limitations

This study has several limitations. Our analysis is restricted by design to thoseIRW datasets with available item text. Although the warehouse's available item text data is growing, the resulting set remains limited. Additional datasets and domains

are needed to assess the generalizability of the findings. The evaluated LLMs are also proprietary and subject to change, which may limit reproducibility. Future studies should include open-source models, document model versions and configurations, and, as the number of models and datasets grows, examine relationships between model characteristics and performance. This includes examining data that may not have been accessed in model training, the lack of which is another unfortunate reality of current LLM paradigms, even among open-weight models.

Additionaly, the hierarchical Bradley–Terry model assumes a particular parametric form for rater discrimination and position bias. Comparing this specification with alternative aggregation models would clarify the robustness of the findings. Finally, we held the prompt design fixed and therefore cannot determine how sensitive LLM judgments are to prompt wording or structure in this specific setting despite evidence CGCoT improves base model performance. Thus, systematically evaluating prompts continues to be an important direction for future work investigating the psychometric properties of LLMs.

7 Conclusion

Paired comparisons elicited from LLMs provide a scalable, nontrivial signal about item difficulty, and hierarchical Bradley–Terry modeling offers a principled way to pool that signal while accounting for differences in rater discrimination and position bias. Nevertheless, high transitivity and strong inter-model agreement did not ensure close correspondence with respondent-calibrated difficulty. LLM consensus is therefore best viewed as evidence for a provisional item ordering, not as psychometric calibration. Respondent data remain necessary when item parameters will support score interpretation or operational assessment decisions.

Acknowledgments

The research reported here was supported by the Institute of Education Sciences, U.S. Department of Education, through Grant R305D240025 to THE LELAND STANFORD JUNIOR UNIVERSITY. The opinions expressed are those of the authors and do not represent views of the Institute or the U.S. Department of Education.

References

Ralph Allan Bradley and Milton E. Terry. 1952. Rank Analysis of Incomplete Block Designs: I. The Method of Paired Comparisons. *Biometrika*, 39(3/4):324–345.

Ulf Böckenholt. 2001. Hierarchical modeling of paired comparison data. *Psychological Methods*, 6(1):49–66.

Xi Chen, Paul N. Bennett, Kevyn Collins-Thompson, and Eric Horvitz. 2013. Pairwise ranking aggregation in a crowdsourced setting. In *Proceedings of the Sixth ACM International Conference on Web Search and Data Mining*, WSDM '13, pages 193–202. Association for Computing Machinery.

Chrzan, Michael L. Pairadigm: A Python Library for Concept-Guided Chain-of-Thought Pairwise Measurement of Scalar Constructs Using Large Language Models.

Matthew R DiGiuseppe and Michael E Flynn. 2026. Scaling Open-Ended Survey Responses Using LLM-Paired Comparisons. *Public Opinion Quarterly*, 90(3):630–656.

Benjamin W. Domingue, Mika Braginsky, Lucy Caffrey-Maffei, Joshua B. Gilbert, Klint Kanopka, Radhika Kapoor, Hansol Lee, Yiqing Liu, Savira Nadela, Guanzhong Pan, Lijin Zhang, Susu Zhang, and Michael C. Frank. 2025. An introduction to the Item Response Warehouse (IRW): A resource for enhancing data usage in psychometrics. *Behavior Research Methods*, 57(10):276.

Mariano Felice and Zeynep Duran Karaoz. 2024. The British Council submission to the BEA 2024 shared task.

Andrew F. Hayes and Klaus Krippendorff. Answering the Call for a Standard Reliability Measure for Coding Data. 1(1):77–89.

Matias Hoyl. Synthetic Student Responses: LLM-Extracted Features for IRT Difficulty Parameter Estimation.

Shengyu Jiang, Chun Wang, and David J. Weiss. 2016. Sample Size Requirements for Estimation of Item Parameters in the Multidimensional Graded Response Model. *Frontiers in Psychology*, 7.

Radhika Kapoor, Sang T. Truong, Nick Haber, Maria Araceli Ruiz-Primo, and Benjamin W. Domingue. Prediction of Item Difficulty for Reading Comprehension Items by Creation of Annotated Item Repository.

James S. Kim, Mary A. Burkhauser, Jackie E. Relyea, Joshua B. Gilbert, Ethan Scherer, Jill Fitzgerald, Douglas Mosher, and Joseph McIntyre. 2023. Replication Data for: A Longitudinal Randomized Trial of a Sustained Content Literacy Intervention from First to Second Grade: Transfer Effects on Students' Reading Comprehension Outcomes.

James S. Kim, Joshua B. Gilbert, Jackie Eunjung Relyea, Patrick Rich, Ethan Scherer, Mary A. Burkhauser, and Johanna N. Tvedt. 2024. Time to transfer: Long-term effects of a sustained and spiraled content literacy intervention in the elementary grades. *Developmental Psychology*, 60(7):1279–1297.

Diana Kolesnikova, Kirill Fedyanin, Abe D. Hofman, Matthieu J. S. Brinkhuis, and Maria Bolsinova. 2026. Estimating Item Difficulty with Large Language Models as Experts.

Ming Li, Han Chen, Yunze Xiao, Jian Chen, Hong Jiao, and Tianyi Zhou. Can LLMs Estimate Student Struggles? Human-AI Difficulty Alignment with Proficiency Simulation for Item Difficulty Prediction. In *Findings of the Association for Computational Linguistics: ACL 2026*, pages 25414–25441. Association for Computational Linguistics.

Hauke Licht, Rupak Sarkar, Patrick Y. Wu, Pranav Goel, Niklas Stoehr, Elliott Ash, and Alexander Miserlis Hoyle. 2025. Measuring Scalar Constructs in Social Science with LLMs. *Preprint*, arXiv:2509.03116.

Mengjie Qian, Guangzhi Sun, Mark Gales, and Kate Knill. 2026. Who can we trust? LLM-as-a-jury for Comparative Assessment.

Rasch, Georg. 1960. *Probabilistic Models for Some Intelligence and Attainment Tests*, volume 1 of *Studies in Mathematical Psychology*. Danmarks Paedagogiske Institut.

Jackie E. Relyea, Joshua B. Gilbert, Mary Burkhauser, Ethan Scherer, Douglas M. Mosher, Zhongyu Wei, Johanna Tvedt, and James S. Kim. 2025. Asset-Based Implementation of Structured Adaptations in an Online Third-Grade Content Literacy Intervention. *Reading Research Quarterly*, 60(4):e70048.

Alexander Scarlatos, Nigel Fernandez, Christopher Ormerod, Susan Lottridge, and Andrew Lan. SMART: Simulated Students Aligned with Item Response Theory for Question Difficulty Prediction. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 25071–25094. Association for Computational Linguistics.

Patrycja Sleboda and Joanna Sokolowska. Measurements of Rationality: Individual Differences in Information Processing, the Transitivity of Preferences and Decision Strategies. 8:1844.

C. Spearman. The proof and measurement of association between two things. 15(1):72–101.

Curtis Tatsuoka. 2002. Data Analytic Methods for Latent Partially Ordered Classification Models. *Journal of the Royal Statistical Society Series C: Applied Statistics*, 51(3):337–350.

Alina A. von Davier, Robert J. Mislevy, and Jiangang Hao. 2022. *Computational Psychometrics: New Methodologies for a New Generation of Digital Learning and Assessment: With Examples in R and Python*. Methodology of Educational Measurement and Assessment. Springer International Publishing.

Howard Wainer, Neil J. Dorans, Ronald Flaugher, Bert F. Green, and Robert J. Mislevy. 2000. *Computerized Adaptive Testing: A Primer*, 2 edition. Routledge.

Xinyi Wang, Hong Jiao, Ming Li, Sydney Peters, Hanna Choi, Tianyi Zhou, and Qingshu Xu. 2026. Can LLMs Really Understand Item Difficulty Levels? Implications for Automated Item Generation Using LLMs. *Preprint*, arXiv:2607.28634.

Patrick Y. Wu, Jonathan Nagler, Joshua A. Tucker, and Solomon Messing. 2024. Concept-Guided Chain-of-Thought Prompting for Pairwise Comparison Scoring of Texts with Large Language Models. In *2024 IEEE International Conference on Big Data (BigData)*, pages 7232–7241.

A CGCoT Prompts

The prompts used were:

- "Summarize the core task, question, behavioral statement, or scenario presented in this item. Item text: {text}".
- "Identify the primary knowledge domain, skill, clinical symptom, attitudinal belief, or psychological construct this item is attempting to measure and give a concise justification. Text: {text}"
- "Analyze the demand or severity of this item. For cognitive/ability items, does it require recall, multi-step reasoning, or synthesis? For psychological/attitudinal items, does endorsing it require a severe level of a symptom, a strongly held belief, or deep introspection? CONCISELY identify any structural or semantic elements—such as complex phrasing, tricky distractors, conditional logic, or strong emotional wording—that increase the cognitive load or the threshold for endorsement. Text: {text}"
- "Based on your analysis, describe the overall level of the underlying latent trait (e.g., high cognitive ability, severe psychological symptomology, strong attitude) required for a respondent to successfully answer or strongly endorse this item. Previous analysis: {text}"

B Hierarchical Bradley-Terry Model Details

B.1 Item Difficulty

The model estimates a single λ_i for each item across all LLMs, allowing us to pool information across models while accounting for model-specific discrimination and position bias. These item strengths are centered to remove translation invariance:

$$\begin{aligned}\lambda_i &= \sigma_\lambda(z_i - \bar{z}) \\ u_k &\sim \mathrm{N}(0,1) \\ \sigma_\lambda &\sim \mathrm{HalfNormal}(1)\end{aligned}$$

Thus $\sum_i \lambda_i = 0$. σ_λ is the standard deviation of the item strengths, z_i is a latent standard-normal random effect, and $\bar{z}$ is the mean of the z_i across all items,

B.2 Rater Position Bias

Another way to conceptualize γ_k is a measure of model k's tendency to prefer the first item *regardless* of item strength. We model γ_k hierarchically, allowing for partial pooling across models and estimating the overall distribution of position bias:

$$\begin{aligned}\gamma_k &= \sigma_\gamma v_k \\ v_k &\sim \mathrm{N}(0,1) \\ \sigma_\gamma &\sim \mathrm{HalfNormal}(0.5)\end{aligned}$$

where σ_γ is the standard deviation of the position bias across models and v_k is a latent standard-normal random effect.

B.3 Rater Discrimination

The discrimination parameter α_k captures the ability of model k to distinguish between items of different strengths. We model α_k hierarchically, allowing for partial pooling across models and estimating the overall distribution of discrimination:

$$\begin{aligned}\log(\alpha_k) &= \sigma_{\log\alpha}(u_k - \bar{u}) \\ u_k &\sim \mathrm{N}(0,1) \\ \sigma_{\log(\alpha)} &\sim \mathrm{HalfNormal}(0.5)\end{aligned}$$

where $\sigma_{\log(\alpha)}$ is the standard deviation of the discrimination across models and u_k is a latent standard-normal random effect.

Centering log discrimination enforces

$$\exp(\frac{1}{K}\sum_k \log(\alpha_k)) = 1$$

so the geometric mean rater discrimination is one. This identifies the global $\alpha - \lambda$ scale without choosing an arbitrary reference model.

C HBT Model Diagnostics

Each dataset-specific HBT model was estimated with four chains, 500 tuning iterations, 1,000 retained draws per chain, and a target acceptance probability of 0.95. Across item strengths, rater effects, and their scale parameters, the largest $\hat{R}$ was 1.01. The minimum bulk and tail effective sample sizes were 594 and 953, respectively. No divergent transitions were observed, and the minimum chain-level Bayesian fraction of missing information was 0.657. Visual inspection of the trace plots for σ_λ and $\sigma_{\log\alpha}$ showed stable mixing with no evident chain separation. Together, these checks provided no indication that the substantive findings were driven by failed convergence.

We evaluated fitted behavior in two complementary ways. First, for every posterior draw we compared the deterministic ordering of λ_i and λ_j with the observed winner, deliberately excluding α_k and γ_k. Mean agreement ranged from 0.716 for `gilbert_meta_104` to 0.871 for `gilbert_meta_103`. Second, we scored each judgment under the complete likelihood. Mean observed-outcome probabilities ranged from 0.653 to 0.828, mean log scores from -0.516 to -0.272, Brier scores from 0.087 to 0.174, and classification accuracy from 0.732 to 0.874. `gilbert_meta_103` had the strongest full-likelihood fit, whereas `frac20` had the lowest predictive accuracy and largest Brier score. For `gilbert_meta_104`, full-likelihood accuracy exceeded deterministic ranking agreement by approximately seven percentage points, indicating that rater discrimination and position bias explained a meaningful portion of its observed choices.

These posterior predictive quantities are in-sample diagnostics: the pairwise judgments being scored are the same judgments used to estimate the model. They measure how well the fitted HBT model represents the observed comparison data, not generalization to unobserved pairs, new items, or new LLM raters.

D Sensitivity Analyses

We examined whether the pooled item ranking depended on the unequal number of models from each provider or on any single LLM.

The provider-balanced model assigned weights so that the three Google models and four OpenAI models contributed equal total likelihood mass. Within each dataset, we compared its posterior-mean item ranks with those from the primary equal-judgment model using Spearman correlation. Correlations ranged from 0.956 for `frac20` to 0.997 for `gilbert_meta_103`; mean absolute rank changes were at most 1.20 positions, and no item moved more than four positions. The main ranking was therefore insensitive to provider composition.

We next refit the model seven times per dataset, excluding one LLM at a time. Of the 56 leave-one-model-out fits, 53 correlated at least 0.944 with the primary ranking, and correlations reached 1.000 in one fit. Across all fits, mean absolute rank changes did not exceed 1.56 positions and the largest change for any item was six positions. The three lower correlations all occurred in the nine-item `gilbert_meta_104` dataset: 0.717 after excluding Gemini 3.7 Flash, 0.767 after excluding Gemini 3 Flash Preview, and 0.917 after excluding Gemini 3.1 Pro Preview. By contrast, the provider-balanced ranking for this dataset remained close to the primary ranking ($\rho = 0.983$). This localized sensitivity is consistent with the greater rank volatility expected in a small item set and indicates that the `gilbert_meta_104` ordering depends more heavily on the Gemini raters. It does not alter the broader finding that no single model determined the pooled rankings in the larger datasets.

All sensitivity fits showed satisfactory convergence for the item-strength parameters: the maximum $\hat{R}$ was 1.01 and the minimum bulk effective sample size was 2,277. Taken together, the provider-balanced and leave-one-model-out analyses support the robustness of the primary ranking, with the explicit exception of model-removal sensitivity in the smallest dataset.

E Additional Tables

The following tables report the complete model-specific validity and transitivity results summarized in the main text.

Model	ρ	p	N
GPT-5.4 Mini	.146	.051	179
GPT-5.4	.195	.009	179
GPT-5.6 Luna	.276	$< .001$	179
GPT-5.6 Sol	.321	$< .001$	179
Gemini 3 Flash	.377	$< .001$	179
Gemini 3.1 Pro	.383	$< .001$	179
Gemini 3.7 Flash	**.402**	$< .001$	**179**

Table 2: Overall model-specific convergent validity across all seven datasets. ρ is the Spearman correlation between each model's standardized Bradley–Terry score and 1PL IRT difficulty; N is the total number of unique items. The strongest correlation is bolded.

Dataset	Model	ρ	p	N
frac20	GPT-5.4 Mini	.448	.048	20
	GPT-5.4	.325	.162	20
	GPT-5.6 Luna	.292	.212	20
	GPT-5.6 Sol	**.644**	**.002**	**20**
	Gemini 3 Flash	.314	.177	20
	Gemini 3.1 Pro	.594	.006	20
	Gemini 3.7 Flash	.350	.130	20
gilbert_meta_2	GPT-5.4 Mini	.365	.113	20
	GPT-5.4	**.435**	**.056**	**20**
	GPT-5.6 Luna	.287	.220	20
	GPT-5.6 Sol	.352	.128	20
	Gemini 3 Flash	.400	.081	20
	Gemini 3.1 Pro	.358	.121	20
	Gemini 3.7 Flash	.388	.091	20
gilbert_meta_7	GPT-5.4 Mini	.012	.945	36
	GPT-5.4	.163	.342	36
	GPT-5.6 Luna	.191	.264	36
	GPT-5.6 Sol	.214	.209	36
	Gemini 3 Flash	**.455**	**.005**	**36**
	Gemini 3.1 Pro	.262	.123	36
	Gemini 3.7 Flash	**.455**	**.005**	**36**
gilbert_meta_8	GPT-5.4 Mini	.254	.183	29
	GPT-5.4	.227	.237	29
	GPT-5.6 Luna	.303	.110	29
	GPT-5.6 Sol	.400	.032	29
	Gemini 3 Flash	**.467**	**.011**	**29**
	Gemini 3.1 Pro	.452	.014	29
	Gemini 3.7 Flash	.465	.011	29
gilbert_meta_102	GPT-5.4 Mini	.127	.461	36
	GPT-5.4	.019	.911	36
	GPT-5.6 Luna	.324	.054	36
	GPT-5.6 Sol	.289	.087	36
	Gemini 3 Flash	.503	.002	36
	Gemini 3.1 Pro	.419	.011	36
	Gemini 3.7 Flash	**.505**	**.002**	**36**
gilbert_meta_103	GPT-5.4 Mini	.243	.203	29
	GPT-5.4	.259	.175	29
	GPT-5.6 Luna	.284	.136	29
	GPT-5.6 Sol	**.464**	**.011**	**29**
	Gemini 3 Flash	.295	.120	29
	Gemini 3.1 Pro	.431	.020	29
	Gemini 3.7 Flash	.367	.050	29
gilbert_meta_104	GPT-5.4 Mini	-.233	.546	9
	GPT-5.4	.600	.088	9
	GPT-5.6 Luna	.083	.831	9
	GPT-5.6 Sol	**.617**	**.077**	**9**
	Gemini 3 Flash	.067	.865	9
	Gemini 3.1 Pro	.567	.112	9
	Gemini 3.7 Flash	.133	.732	9

Table 3: Model-specific convergent validity by dataset. ρ is the Spearman correlation between each model's standardized Bradley–Terry score and 1PL IRT difficulty; N is the number of unique items. The strongest correlation within each dataset is bolded; the two Gemini Flash models tie on gilbert_meta_7.

Dataset	Model	Transitivity (%)	Triplets
frac20	GPT-5.4 Mini	99.4	641
	GPT-5.4	98.9	641
	GPT-5.6 Luna	95.9	641
	GPT-5.6 Sol	98.3	641
	Gemini 3 Flash	99.4	641
	Gemini 3.1 Pro	98.4	641
	Gemini 3.7 Flash	99.5	641
gilbert_meta_2	GPT-5.4 Mini	87.9	686
	GPT-5.4	98.3	686
	GPT-5.6 Luna	99.0	686
	GPT-5.6 Sol	99.7	686
	Gemini 3 Flash	99.1	686
	Gemini 3.1 Pro	100.0	686
	Gemini 3.7 Flash	99.6	686
gilbert_meta_7	GPT-5.4 Mini	91.4	706
	GPT-5.4	98.2	706
	GPT-5.6 Luna	97.7	706
	GPT-5.6 Sol	99.6	706
	Gemini 3 Flash	98.4	706
	Gemini 3.1 Pro	99.7	706
	Gemini 3.7 Flash	99.7	706
gilbert_meta_8	GPT-5.4 Mini	96.8	686
	GPT-5.4	99.3	686
	GPT-5.6 Luna	99.4	686
	GPT-5.6 Sol	99.4	686
	Gemini 3 Flash	95.2	686
	Gemini 3.1 Pro	99.9	686
	Gemini 3.7 Flash	99.6	686
gilbert_meta_102	GPT-5.4 Mini	91.9	755
	GPT-5.4	96.3	755
	GPT-5.6 Luna	99.6	755
	GPT-5.6 Sol	99.7	755
	Gemini 3 Flash	98.8	755
	Gemini 3.1 Pro	100.0	755
	Gemini 3.7 Flash	99.5	755
gilbert_meta_103	GPT-5.4 Mini	95.5	719
	GPT-5.4	99.7	719
	GPT-5.6 Luna	99.7	719
	GPT-5.6 Sol	99.7	719
	Gemini 3 Flash	98.1	719
	Gemini 3.1 Pro	99.3	719
	Gemini 3.7 Flash	99.4	719
gilbert_meta_104	GPT-5.4 Mini	76.2	84
	GPT-5.4	85.7	84
	GPT-5.6 Luna	96.4	84
	GPT-5.6 Sol	98.8	84
	Gemini 3 Flash	100.0	84
	Gemini 3.1 Pro	100.0	84
	Gemini 3.7 Flash	100.0	84

Table 4: Observed triplet transitivity by model and dataset. Transitivity is the percentage of evaluated item triplets whose pairwise judgments formed a consistent ordering. Triplets gives the number evaluated for each model.

Implementing Generative AI Assessment Guidelines in a Mexican University

Melchor Sánchez-Mendiola, Enrique Buzo-Casanova, Elibidú Ortega-Sánchez, Manuel García-Minjares, and Adriana Cruz Juárez
Coordination of Evaluation, Innovation and Educational Development (CEIDE)
National Autonomous University of Mexico (UNAM), Mexico City, Mexico
Correspondence: melchorsm@unam.mx

Abstract

Generative artificial intelligence (GenAI) complicates authorship and the interpretation of evidence of learning, yet institutional responses often remain limited to prohibition, detection, or broad principles. This work-in-progress case study examines how the National Autonomous University of Mexico (UNAM) developed and launched three coordinated guidelines for GenAI use in assessment across upper-secondary, undergraduate, and graduate education. The corpus comprised 120 pages of guidelines and project records. Structured document analysis, process tracing, and framework mapping identified shared governance principles, level-specific adaptations, and a pathway from institutional guidance to course-level assessment decisions. The documents share a learning-centered core of validity, fairness, transparency, academic integrity, privacy, bias awareness, metacognition, and human oversight. Their emphasis differs by level: gradual, guided use in upper-secondary education; assessment-cycle redesign and rubrics at undergraduate level; and AI literacy, research integrity, curated-source tools, and oral dialogue at graduate level. Development occurred through five collaborative phases over approximately 11 months. The documents received 5,276 downloads during their first month, indicating early reach rather than adoption or effectiveness. We propose a coherence-with-adaptation model and an implementation evaluation plan addressing acceptability, appropriateness, feasibility, adoption, fidelity/adaptation, reach, and sustainability.

1 Introduction

GenAI can draft and revise text, synthesize information, translate, create images and code, simulate dialogue, generate assessment tasks, and provide feedback. These capabilities make polished products increasingly ambiguous as evidence of unaided learning. A written submission may represent independent reasoning, legitimate assistance, close human-AI collaboration, or substantial substitution of the learner's work. The validity question is therefore not simply whether GenAI was used, but whether the evidence supports the intended interpretation of learning under the assessment conditions specified (American Educational Research Association et al., 2014).

Institutional responses initially emphasized misconduct and detection. Those concerns remain important, but prohibition-only approaches can obscure legitimate educational uses, inequitable access, privacy risks, and preparation for AI-mediated professional work (Dabis & Csáki, 2024; UNESCO, 2023). Reviews of university guidance show a shift toward course-level communication, assessment redesign, responsible use, and stakeholder support (An et al., 2025; Moorhouse et al., 2023). Effective adoption also requires clear roles, infrastructure, implementation supports, and continuing revision rather than a static policy document (Chan, 2023; Jin et al., 2025; Robert & McCormack, 2024).

Assessment governance must protect tasks in which unaided performance is part of the construct while permitting or requiring GenAI when critical use of the technology is itself a learning outcome. It must communicate a common institutional position without erasing differences among disciplines, educational levels, assessment purposes, and faculty judgment. It must also translate broad principles into usable assignment instructions, rubrics, disclosure rules, and evidence-collection strategies.

This paper examines three guidelines for GenAI use in educational assessment at the upper-secondary, undergraduate, and graduate levels in a public university in Mexico (Universidad Nacional Autónoma de México [UNAM], 2026a, 2026b, 2026c). We ask: (1) How can guidance remain coherent across levels while preserving

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 37–44
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

developmental and disciplinary relevance? (2) What process supported collaborative guideline development? (3) What implications follow for assessment validity and implementation?

2 Institutional Context and Development Process

UNAM is a large public university spanning upper-secondary, undergraduate, and graduate education across heterogeneous disciplines and organizational settings. Assessment decisions are distributed among academic bodies, programs, courses, and faculty. Institutional survey work before publication also indicated uneven GenAI uptake and limited assessment use in some sectors, including graduate education (Benavides-Lara et al., 2025). A single prescriptive rule would therefore be difficult to enact and pedagogically undesirable. The project pursued coordinated guidance rather than a uniform binding regulation.

The university Educational Evaluation Council and the Coordination of Evaluation, Innovation and Educational Development (CEIDE) convened three collaborative networks representing the three educational levels. Their shared objective was to provide practical, ethical, and pedagogically grounded guidance for assessment of, for, and as learning. The documents are advisory, they complement existing norms and faculty decisions rather than replacing them. All were published in Spanish under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 license (UNAM, 2026a, 2026b, 2026c).

Project records document five phases. Start-up and agreements established the objective, structure, references, writing conventions, schedule, and shared workspace. Content development used approximately biweekly meetings and distributed writing. Internal integration and review merged contributions, checked coherence and coverage, and standardized references. Validation and final revision incorporated coordinator review, stylistic editing, and feedback. Closure and delivery included approval, production, and public release. The process lasted approximately 11 months.

The networks worked from a common architecture but adapted examples and emphasis. This operationalized the project's central design principle: coherence need not mean uniformity. A shared core defined values and decision requirements; each guide translated them into developmentally and disciplinarily appropriate recommendations.

3 Methods

3.1 Design and Corpus

We conducted an implementation case study using document analysis, process tracing, and framework mapping. The primary corpus consisted of the final guidelines: 32 pages for upper-secondary, 57 for undergraduate, and 31 for graduate education, totaling 120 pages. Supporting materials included the collaborative work plan, network presentations, integrated review materials, and the April 2026 launch presentation. No student records, faculty performance data, or individual-level survey responses were analyzed. Table 1 summarizes the corpus and the orientation of each guide.

Table 1: Documentary corpus and level-specific orientation of the three UNAM GenAI assessment guidelines.

Level and document	Scope and audience	Distinctive emphasis	Representative operational guidance
Upper-secondary (32 pages)	Faculty, academic staff, assessment designers, and students; developmental context includes many learners below legal adulthood.	Gradual, guided use; verification; metacognition; privacy and equitable access.	Guided self-assessment and generative quizzes in early semesters; cases and simulations in intermediate stages; integrative projects, source checking, use declarations, and reflective learning records in later stages.
Undergraduate (57 pages)	Diverse degree programs, disciplines, course formats, and faculty assessment practices.	Full assessment-cycle redesign from planning through metaevaluation; explicit local rules and rubrics.	Specify AI-use conditions; diversify evidence; communicate criteria; retain process evidence where relevant; use rubrics for critical and creative GenAI use; review validity, bias, and assessment consequences.
Graduate (31 pages)	Specialization, master's, and doctoral education across 42 programs and research-intensive contexts.	AI literacy, prompt quality, research integrity, disciplinary application, and authorship verification.	Combine no-AI, AI-assisted, and AI-integrated tasks; use curated-source/RAG environments; require transparent declaration; use dialogue or oral defense when undeclared use is suspected.

3.2 Document Analysis and Framework Mapping

The unit of analysis was a substantive principle, recommendation, warning, example, or implementation support. A comparison matrix used seven a priori domains from assessment and AI-governance literature: learning purpose, validity, fairness and equity, transparency, academic integrity, privacy and data protection, and human oversight. Metacognition, AI literacy, and implementation support emerged during review. The analysis was interpretive rather than frequency based; it sought functional similarities and meaningful adaptations, not topic counts.

Recommendations were then mapped to two complementary frameworks. The first distinguished assessment of, for, and as learning (Black & Wiliam, 2009; Sánchez Mendiola & Martínez González, 2024). The second distinguished no-AI, AI-assisted, and AI-integrated assessment (Digital Education Council & Pearson, 2025). The relevant condition depends on the construct: no AI when independent performance is the target; assistance when the tool supports but does not replace reasoning; and integration when responsible, critical GenAI use is part of the intended competence.

Process tracing reconstructed the development phases and governance tensions. Written products were compared with project records to distinguish common commitments from level-specific adaptations.

3.3 Early Reception and Planned Implementation Evaluation

Available reception evidence comprised the public launch and aggregate download counts. Downloads were interpreted only as reach and interest; they do not show that faculty read, accepted, adopted, or implemented the guidance, or that assessment quality improved.

For the next phase, we adapted implementation outcomes proposed by Proctor et al. (2011), informed by Nilsen's (2015) distinction among process, determinant, and evaluation frameworks. Planned outcomes are acceptability, appropriateness, feasibility, adoption, fidelity and local adaptation, reach, and sustainability. This distinction matters because dissemination can be extensive without meaningful practice change.

4 Results

4.1 A Common Assessment-Governance Core

Across the documents, GenAI is framed as pedagogical augmentation rather than a substitute for learner agency or professional judgment. The shared core combines three functions. First, it protects validity: learning purpose and construct should precede technology choice; task conditions, evidence, and interpretation should align; and AI-generated instruments or feedback require human review. Second, it establishes ethical and procedural safeguards: equitable access, advance communication of permitted uses, disclosure of material AI assistance, protection of sensitive data, verification of claims, bias awareness, and human responsibility for consequential judgments. Third, it treats metacognition and AI literacy as educational aims: learners should explain how they used a tool, what they accepted or rejected, and how they checked the result.

Table 2 translates these commitments into operational safeguards. Collectively, they shift the governing question from "Is AI allowed?" to "What use condition produces defensible evidence for this learning purpose?" The documents therefore combine protective and developmental aims: preventing substitution, bias, privacy loss, and invalid inference while supporting feedback, self-regulation, critical evaluation, and professional preparation.

Table 2: Common assessment-governance principles identified across the three guidelines.

Governance principle	Operational translation	Primary safeguard
Learning centrality and validity	Define the construct, learning outcome, task, and evidence before selecting an AI tool.	Prevents technology-driven design and construct drift.
Explicit AI-use condition	State whether AI is prohibited, permitted as assistance, or integrated as part of the competence.	Clarifies administration conditions and score interpretation.
Transparency and integrity	Communicate rules in advance and require disclosure of material AI assistance.	Supports authorship, accountability, and consistent expectations.
Fairness and access	Consider devices, connectivity, accessibility, AI literacy, and non-AI alternatives.	Limits opportunity gaps and exclusion.

Governance principle	Operational translation	Primary safeguard
Privacy and data protection	Avoid sensitive inputs; anonymize information and prefer institutionally approved tools.	Protects learners, faculty, and confidential academic material.
Verification and bias awareness	Corroborate claims, sources, and calculations; examine cultural and algorithmic bias.	Reduces error, confabulation, and inequitable consequences.
Human oversight	Retain faculty responsibility for consequential interpretation, feedback, and adjudication.	Preserves professional accountability and contextual judgment.
Metacognition and AI literacy	Ask learners to explain prompts, choices, corrections, rejected outputs, and limitations.	Makes critical human contribution and responsible use visible.

4.2 Level-Specific Adaptation

The upper-secondary guide emphasizes developmental readiness and guided progression. Because many learners are below legal adulthood, privacy and protection receive particular attention. It recommends gradual, guided, and pedagogically intentional use. Early applications include guided self-assessment, generative quizzes, and concept maps; later stages introduce cases, simulations, interdisciplinary problems, integrative projects, and more independent critique. Student guidance emphasizes source verification, citation, privacy, reflection on the system's contribution, and learning records. Rubrics may assess recognition of AI limitations, verification of information, and explanation of how the tool supported learning (UNAM, 2026a).

The undergraduate guide treats assessment as a full decision cycle: planning, evidence collection, interpretation, feedback, decision making, and metaevaluation. This prevents AI from being reduced to item generation or misconduct detection and accommodates examinations, projects, portfolios, performance assessments, self-assessment, and peer assessment. Faculty are advised to specify the AI-use condition for each task, communicate criteria in advance, preserve process evidence when relevant, and use rubrics that evaluate critical and creative AI use rather than polished output alone. Metaevaluation asks whether instruments were valid and reliable, criteria clear, bias reduced, and deep learning and justice promoted (UNAM, 2026b).

The graduate guide addresses specialization, master's, and doctoral education across 42 programs. It emphasizes AI literacy, prompt quality, disciplinary application, research integrity, and preservation of higher-order analytical and creative work. It combines no-AI tasks for foundational competence, AI-assisted work in which students remain responsible for validation and synthesis, and AI-integrated tasks that assess critical and ethical use. For evidence-intensive activities, it recommends retrieval-augmented or curated-source environments grounded in faculty-selected scientific documents. AI use should be declared explicitly. When undeclared use is suspected, dialogue and an oral response or defense are recommended before reliance on AI-detection software, treating authorship verification as an educational and evidentiary problem (UNAM, 2026c).

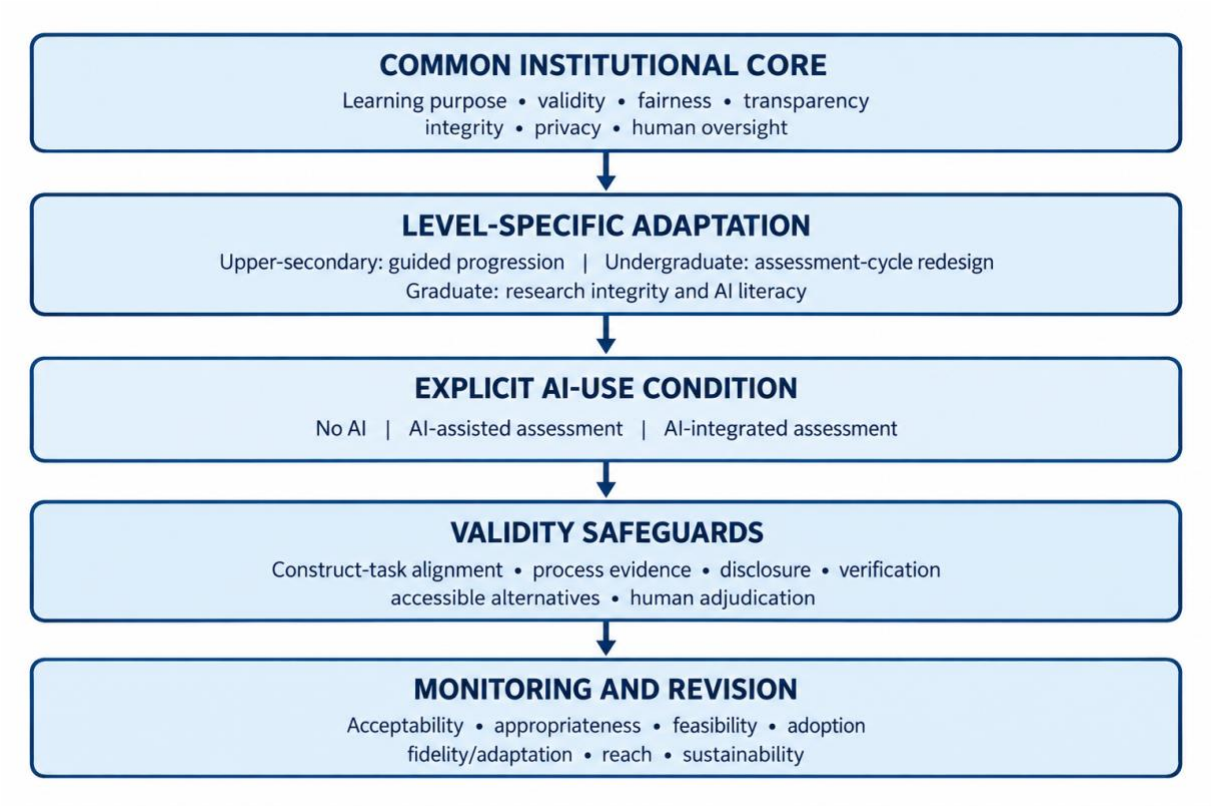
4.3 From Principles to Course-Level Assessment Decisions

Cross-document comparison yielded a five-layer policy-to-practice pathway (Figure 1): shared institutional core; level and disciplinary adaptation; explicit task-level AI-use condition; validity safeguards; and monitoring and revision. The safeguards include construct-task alignment, proportionate process evidence, disclosure, source verification, human adjudication, and accessible alternatives.

The pathway places tool rules downstream from construct decisions. An unaided explanation may support an inference about independent understanding; responsible GenAI-assisted synthesis requires tool access and criteria for verification, attribution, and judgment. Neither condition is inherently more authentic; each is defensible only in relation to the intended claim. Similarly, drafts, prompt logs, annotations, oral explanations, or reflective statements are useful only when they illuminate construct-relevant reasoning without disproportionate burden or privacy risk.

At minimum, an assessment that involves GenAI should communicate the learning outcome, AI-use condition, work that must remain human, disclosure requirement, possible process evidence, verification expectations, access and privacy safeguards, and locus of final judgment.

Figure 1: Coherence-with-adaptation pathway from institutional principles to task-level assessment conditions and implementation monitoring.



4.4 Early Reach and Current Implementation Status

The guidelines were launched publicly on April 29, 2026 and released as open-access documents. During their first month, they accumulated 5,276 downloads, indicating substantial early visibility for Spanish-language, assessment-specific GenAI guidance. The metric does not show full reading, perceived usefulness, adoption, or practice change. Current evidence therefore supports conclusions about development, documentary coherence, level-specific adaptation, and dissemination reach, not implementation fidelity, effectiveness, or learner outcomes.

5 Discussion

5.1 Coherence Without Pedagogical Uniformity

The main contribution is a governance design that combines a stable institutional core with differentiated operational guidance. Existing analyses repeatedly identify the need for clear course-level communication, ethical guidance, assessment redesign, and stakeholder responsibilities (An et al., 2025; Chan, 2023; Jin et al., 2025). The UNAM case adds a cross-level perspective: upper-secondary, undergraduate, and graduate education share values but differ in developmental readiness, disciplinary complexity, research intensity, and learner autonomy. One undifferentiated document could obscure these differences; unrelated documents could create inconsistency.

The distinction between policy and guidance is consequential. Formal rules can establish mandatory boundaries for privacy, integrity, or high-stakes decisions, while flexible guidance helps faculty interpret those boundaries for particular constructs and tasks (Robert & McCormack, 2024). The documents primarily perform the second role. Their adaptability is a strength, but optional guidance may be enacted unevenly; publication must be followed by implementation support.

5.2 Implications for Measurement and Validity

GenAI makes the assessment condition part of the evidentiary argument. A product created under unspecified conditions has reduced interpretability because the contributions of learner, tool, and other sources are unclear. Explicit AI-use conditions improve interpretive clarity, but validity still

depends on alignment among construct, task, scoring criteria, and consequences.

Three implications follow. Institutions should preserve selected no-AI assessments when independent recall, reasoning, communication, or performance is essential. When assistance is permitted, tasks should elicit verification, selection, correction, and explanation rather than reward unexamined output fluency. When AI use is part of the construct, criteria should assess problem framing, source judgment, bias awareness, transparency, and accountability, not prompt mechanics alone.

Richer evidence systems should not imply surveillance of every interaction. Excessive monitoring can create privacy risks and construct-irrelevant burden. A defensible approach samples the minimum additional evidence needed: selected drafts, a short use declaration, annotated sources, an oral explanation, or a reflective account. The measurement contribution is therefore a decision framework for specifying when and why additional evidence is warranted, with human judgment retained for contextual interpretation.

5.3 From Dissemination to Implementation

It is important to emphasize that institutional reach is not implementation. Exposure must be followed by evidence that intended users find the guidance acceptable and appropriate, can feasibly apply it, and incorporate it into assessment design. Faculty development is central: educators need support to define AI-use conditions, redesign tasks, construct criteria for critical AI use, verify generated material, protect data, and respond fairly to suspected misuse.

The next phase should combine surveys and interviews with review of syllabi, assignment instructions, rubrics, and assessment plans. Disciplinary case studies can examine how faculty redesign tasks and how students experience the expectations. Program-level monitoring can assess reach and unequal access to training or approved tools. Recurring questions and implementation problems should feed revision, while successful adaptations can become reusable templates. This feedback loop treats the guidelines as living governance infrastructure rather than a finished publication.

5.4 Transferability

The design problem extends beyond UNAM. Multi-level, multi-campus, or professionally diverse institutions often need common expectations without erasing local expertise. The coherence-with-adaptation model may be especially useful where faculty autonomy is strong and central guidance must coexist with disciplinary variation.

The case also broadens the geographical and linguistic base of GenAI governance work, much of which has examined English-language policies from highly ranked institutions in high-income settings (An et al., 2025; Moorhouse et al., 2023). Spanish-language, open-access guidance from a large Latin American public university foregrounds digital inequality, developmental protection, and local faculty capacity. Transfer requires adaptation rather than replication, but the sequence (common principles, contextual translation, task-level rules, and implementation evaluation) is broadly applicable.

6 Limitations and Next Steps

The authors participated in developing or coordinating the guidelines, providing process access but creating a risk of favorable interpretation. We therefore distinguish documentary findings from impact claims and make institutional positionality explicit. Future evaluation should include non-participant users and, where feasible, independent analysis.

The study contains no faculty or student experience data, observed adoption, fidelity, assessment-quality ratings, or learning outcomes; downloads cannot establish use. The documents also differ in length, style, and operational detail because distinct networks produced them, and the interpretive analysis does not estimate theme prevalence. Finally, rapidly changing technologies and regulations may make tool-specific recommendations obsolete even when governing principles remain relevant.

The planned mixed-method evaluation will operationalize seven implementation outcomes (Table 3). Where redesigned assessments are implemented, later studies should also examine validity evidence, fairness, faculty workload, student learning processes, and unintended consequences.

Table 3: Planned evaluation of implementation outcomes, adapted from Proctor et al. (2011).

Implementation outcome	Key question	Candidate evidence for the next phase
Acceptability	Do intended users consider the guidance clear and useful?	Faculty and student surveys; interviews; recurring questions.
Appropriateness	Does it fit the educational level, discipline, and assessment purpose?	Ratings by level/program; disciplinary case studies; expert review.
Feasibility	Can users apply it with available time, expertise, access, and infrastructure?	Workload estimates; barriers/enablers; privacy and access audits.
Adoption	Are recommendations used in real assessment practice?	Syllabus statements; assignment instructions; rubrics; program guidance.
Fidelity and adaptation	Which core principles are retained, and how is guidance modified locally?	Document comparison; faculty narratives; assessment-design audits.
Reach	Which units and user groups encounter and use the resources?	Downloads, event participation, training coverage, unit-level distribution.
Sustainability	Are responsibilities, updates, support, and version control maintained?	Review schedule; governance ownership; updated editions; continuing faculty development.

7 Conclusion

The UNAM case demonstrates an assessment-governance approach that moves beyond a binary choice between permitting and prohibiting GenAI. Its central principle is coherence with adaptation: a common learning-centered and ethical core translated for upper-secondary, undergraduate, and graduate education. The policy-to-practice pathway begins with the construct, specifies the AI-use condition, identifies proportionate evidence and safeguards, and ends with monitoring and revision.

The project remains a work in progress. Published guidance and strong early reach demonstrate institutional capacity to communicate a framework, not practice change or improved assessment. The next test is whether faculty and programs use it to produce more interpretable, fair, transparent, and educationally meaningful evidence of learning. In the GenAI era, institutions must specify the human reasoning, judgment, and responsibility their assessments are intended to reveal.

Ethical Considerations and Researcher Positionality

This phase analyzed public institutional documents, project records, and aggregate download counts. It involved no individual-level student or faculty data. All authors participated in the initiative, and this dual role is acknowledged as a potential source of interpretive bias.

Acknowledgments

The authors thank the Educational Evaluation Council; the upper-secondary, undergraduate, and graduate collaborative networks; the faculty and academic staff who drafted and reviewed the guidelines; and CEIDE teams supporting editing, design, publication, and dissemination.

References

American Educational Research Association, American Psychological Association, & National Council on Measurement in Education. (2014). Standards for educational and psychological testing. American Educational Research Association. https://www.testingstandards.net/open-access-files.html

An, Y., Yu, J. H., & James, S. (2025). Investigating the higher education institutions' guidelines and policies regarding the use of generative AI in teaching, learning, research, and administration. International Journal of Educational Technology in Higher Education, 22, Article 10. https://doi.org/10.1186/s41239-025-00507-3

Benavides-Lara, M. A., Rendón Cazales, V. J., Escalante Rivas, N., Martínez Hernández, A. M. del P., & Sánchez Mendiola, M. (2025). Presencia y uso de la inteligencia artificial generativa en la Universidad Nacional Autónoma de México. Revista Digital Universitaria, 26(1). https://doi.org/10.22201/ceide.16076079e.2025.26.1.10

Black, P., & Wiliam, D. (2009). Developing the theory of formative assessment. Educational Assessment, Evaluation and Accountability, 21, 5-31. https://doi.org/10.1007/s11092-008-9068-5

Chan, C. K. Y. (2023). A comprehensive AI policy education framework for university teaching and learning. International Journal of Educational Technology in Higher Education, 20, Article 38. https://doi.org/10.1186/s41239-023-00408-3

Dabis, A., & Csáki, C. (2024). AI and ethics: Investigating the first policy responses of higher education institutions to the challenge of generative AI. Humanities and Social Sciences Communications, 11, Article 1006. https://doi.org/10.1057/s41599-024-03526-z

Digital Education Council & Pearson. (2025). The next era of assessment: A global review of AI in assessment design. https://www.digitaleducationcouncil.com/post/the-next-era-of-assessment-a-global-review-of-ai-in-assessment-design

Jin, Y., Yan, L., Echeverria, V., Gašević, D., & Martinez-Maldonado, R. (2025). Generative AI in higher education: A global perspective of institutional adoption policies and guidelines. Computers and Education: Artificial Intelligence, 8, Article 100348. https://doi.org/10.1016/j.caeai.2024.100348

Moorhouse, B. L., Yeo, M. A., & Wan, Y. W. (2023). Generative AI tools and assessment: Guidelines of the world's top-ranking universities. Computers and Education Open, 5, Article 100151. https://doi.org/10.1016/j.caeo.2023.100151

Nilsen, P. (2015). Making sense of implementation theories, models and frameworks. Implementation Science, 10, Article 53. https://doi.org/10.1186/s13012-015-0242-0

Proctor, E., Silmere, H., Raghavan, R., Hovmand, P., Aarons, G., Bunger, A., Griffey, R., & Hensley, M. (2011). Outcomes for implementation research: Conceptual distinctions, measurement challenges, and research agenda. Administration and Policy in Mental Health and Mental Health Services Research, 38(2), 65-76. https://doi.org/10.1007/s10488-010-0319-7

Robert, J., & McCormack, M. (2024). 2024 EDUCAUSE action plan: AI policies and guidelines. EDUCAUSE. https://library.educause.edu/resources/2024/5/2024-educause-action-plan-ai-policies-and-guidelines

Sánchez Mendiola, M., & Martínez González, A. (Eds.). (2024). Evaluación y aprendizaje en educación universitaria: Estrategias e instrumentos (2.ª reimp.). Universidad Nacional Autónoma de México, Coordinación de Evaluación, Innovación y Desarrollo Educativos. https://www.ceide.unam.mx/index.php/evaluacion-y-aprendizaje-en-educacion-universitaria/

UNESCO. (2023). Guidance for generative AI in education and research. https://www.unesco.org/en/articles/guidance-generative-ai-education-and-research

Universidad Nacional Autónoma de México. (2026a). Guía de uso de inteligencia artificial generativa en la evaluación educativa en bachillerato [Guidelines for the use of generative artificial intelligence in educational assessment at the upper-secondary level]. Consejo de Evaluación Educativa; Red Colaborativa del Bachillerato; Coordinación de Evaluación, Innovación y Desarrollo Educativos. https://www.ceide.unam.mx/index.php/informes/#CEE

Universidad Nacional Autónoma de México. (2026b). Guía de uso de inteligencia artificial generativa en la evaluación educativa en licenciatura [Guidelines for the use of generative artificial intelligence in undergraduate educational assessment]. Consejo de Evaluación Educativa; Red Colaborativa de Licenciatura; Coordinación de Evaluación, Innovación y Desarrollo Educativos. https://www.ceide.unam.mx/index.php/informes/#CEE

Universidad Nacional Autónoma de México. (2026c). Guía de uso de inteligencia artificial generativa en la evaluación educativa en posgrado [Guidelines for the use of generative artificial intelligence in graduate educational assessment]. Consejo de Evaluación Educativa; Red Colaborativa de Posgrado; Coordinación de Evaluación, Innovación y Desarrollo Educativos. https://www.ceide.unam.mx/index.php/informes/#CEE

TRACE: Automated Multi-Granularity Analysis of Text Revision

Yu Tian, Katerina Christhilf, Andrew Potter, Jessica Early
Steve Graham, and Danielle McNamara
Arizona State University
{ytian126,kchristh,ahpotter,Jessica.Early,steve.graham,dsmcnama}@asu.edu

Abstract

This paper introduces TRACE, an automated framework for analyzing textual revisions across drafts. TRACE aligns initial and revised texts and infers interpretable revision operations, such as insert, delete, replace, and move. TRACE supports AI-powered revision analytics, writing assessment, automated feedback, and studies of human–AI writing collaboration.

1 Introduction

Revision is a central mechanism through which writers refine meaning, reorganize discourse, and improve communicative effectiveness (Faigley and Witte, 1981). In the age of generative AI, revision has become even more consequential. Writers may now compose through repeated interaction with AI, revise AI-generated drafts, or selectively integrate AI-generated suggestions into their own texts (Cengiz et al., 2025; Koltovskaia et al., 2024). As a result, revision can provide evidence of writing effort, strategy use, and patterns of human-AI collaboration.

However, large-scale revision analysis remains methodologically challenging. Manual coding can capture the complexity of revision, but it is difficult to scale (e.g., Gao et al., 2019; Min, 2006). Keystroke logging tools provide detailed temporal evidence about revision processes, but they require specialized data collection environments and cannot be applied retrospectively to existing draft pairs (Leijten and Van Waes, 2013; Tian and Cushing, 2025). Computational approaches provide useful foundations for revision analysis (e.g., Hashemi and Schunn, 2014). Prior work employed advanced methods for low-level edit extraction, sentence alignment, and span-level comparison. However, existing approaches do not fully address the need for a unified, interpretable representation of revision operations across multiple levels of granularity. Low-level diff methods can identify fine-grained textual changes, but they may fragment meaningful revisions. Sentence- and span-alignment methods improve interpretability, but they often focus on correspondence or edit extraction rather than explicitly modeling operation categories such as preservation, replacement, and movement within a single framework. This creates a need for automated methods that can estimate interpretable revision operations from source-target draft pairs while preserving both structural and local textual information.

This paper introduces TRACE, a deterministic framework for Text Revision Analysis via Change Estimation. TRACE focuses on an initial step in revision analysis: estimating interpretable revision operations from source-target draft pairs. It integrates span-level matching, operation inference, and token-level refinement into a unified representation of textual change. TRACE distinguishes preserved text, insertions, deletions, replacements, and moved spans, while retaining links to the original source and target texts. This representation can provide a foundation for subsequent human annotation, scoring, and automated modeling of revision purpose, depth, and quality.

This paper makes three contributions. First, it introduces TRACE as a deterministic framework for estimating interpretable revision operations from source-target draft pairs. Second, it provides a mixed-granularity representation that combines span-level operation inference with token-level refinement and explicitly models moved text. Third, it validates TRACE against process-level evidence from keystroke logging, showing that revision operations inferred retrospectively from draft pairs correspond to observed writing-process actions.

2 Related Work

Automatic comparison of document versions has traditionally relied on sequence alignment and edit-

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 45–51
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

distance algorithms. Levenshtein's edit distance formalized textual difference as the minimum number of elementary operations needed to transform one sequence into another, typically through insertions, deletions, and substitutions (Lcvenshtcin, 1966). Wagner and Fischer (1974) further established edit distance as a dynamic-programming framework for string-to-string correction. Myers (1986) later introduced an efficient imagedifference algorithm for computing shortest edit scripts, which became foundational for practical file comparison and version-control systems. These methods are deterministic and computationally efficient, but they optimize edit cost rather than revision interpretability. As a result, meaningful writing operations, such as sentence movement or partial rewriting, may be represented as fragmented low-level edits.

NLP research has addressed this limitation by aligning larger linguistic units and modeling revision intentions. Sentence alignment has been used to identify corresponding units across related documents, even when wording differs (Barzilay and Elhadad, 2003). In educational writing research, Zhang and Litman (2014) provides a foundational point of comparison. Their approach segments drafts into sentences, aligns sentences using semantic similarity and global contextual information, and infers sentence-level operations such as addition, deletion, and modification. Their work demonstrated that sentence-level alignment can support interpretable revision analysis and downstream classification of revision purposes in argumentative writing.

More recent work has extended revision analysis toward finer-grained and domain-specific representations. For example, Jiang et al. (2022) introduced arXivEdits, a corpus and computational framework for studying scientific writing revision across full-paper versions. Their automated pipeline first aligns sentences across document versions and then formulates edit extraction as a span-alignment problem within aligned sentence pairs. This approach produces fine-grained and explainable edits compared with standard diff-based comparison, demonstrating the value of combining sentence-level correspondence with span-level edit extraction. Together, these research endeavors illustrate a broader shift from low-level differencing toward more interpretable representations of document change.

TRACE builds on this line of work but differs in how revision operations are estimated and represented. Whereas Zhang and Litman (2014) supports sentence-level revision extraction, and Jiang et al. (2022) extracts fine-grained span edits within scientific writing revisions, TRACE focuses on producing a unified operation-level representation across multiple granularities. In addition, prior revision-analysis studies have typically relied on manual annotation, sentence-alignment evaluation, or naturally occurring revision histories, such as Wikipedia edit histories (Bronner and Monz, 2012; Daxenberger and Gurevych, 2013). In contrast, TRACE is evaluated against process-level evidence from keystroke logging. This validation tests whether revision operations inferred retrospectively from source-target draft pairs correspond to actual revision actions observed during writing. Thus, TRACE's main contribution is an interpretable, multi-granularity revision-estimation framework that explicitly represents text movement and is empirically grounded through comparison with keystroke-log evidence.

3 TRACE Pipeline

3.1 Problem Definition

Let T^{src} denote a source draft and T^{tgt} denote its revised version. The task is to estimate the textual changes that transform T^{src} into T^{tgt}. However, this task requires more than identifying a minimal sequence of character-level edits. Character-level diff methods prioritize a globally minimal edit path, which can fragment meaningful revisions such as sentence splitting, sentence merging, replacement, and movement. As a result, structural interpretability may be obscured.

At the same time, approaches based primarily on semantic similarity may introduce a different problem. Although semantic similarity can help identify related textual units, it may also align sentences that are conceptually similar but structurally different. This is especially important when text movement is an explicit operation category, because semantically similar but newly written or substantially rewritten text may otherwise be misidentified as relocated text.

Given these challenges, TRACE aims to transform a source-target draft pair into an ordered, interpretable representation of revision behavior. The goal is to identify how textual spans in T^{src} are preserved, deleted, inserted, replaced, or moved in T^{tgt}, while preserving character-level links to the original texts.

TRACE represents revision operations using the set

$$\mathcal{O} = \{\textit{preserve}, \textit{insert}, \textit{delete}, \textit{replace}, \textit{move}\}.$$

3.2 Step 1: Sentence- and Span-Level Matching

TRACE first identifies corresponding and non-corresponding spans between T^{src} and T^{tgt}. Both drafts are segmented into sentence sequences:

$$\mathbf{S} = (S_1, \ldots, S_M), \qquad \mathbf{T} = (T_1, \ldots, T_N),$$

where M and N are the numbers of sentences in the source and target drafts.

Very short sentences are merged with adjacent sentences to reduce unstable matches caused by short fragments. In the current implementation, sentences with fewer than $\tau = 4$ tokens are merged, preferably with the following sentence. If the short sentence is final, it is merged with the preceding sentence.

For each source-target candidate pair, TRACE computes token-level normalized edit-distance similarity. Let a and b denote token sequences from candidate source and target spans, and let $d(a, b)$ denote their token-level Levenshtein distance. The normalized distance and similarity are defined as:

$$d_{\text{norm}}(a, b) = \frac{d(a, b)}{\max(|a|, |b|)},$$
$$\text{sim}(a, b) = 1 - d_{\text{norm}}(a, b).$$

For each source sentence S_i, where $i \in \{1, \ldots, M\}$, TRACE first selects the best-matching target sentence:

$$\hat{\jmath}(i) = \arg \max_{j \in \{1, \ldots, N\}} \text{sim}(S_i, T_j).$$

This produces an initial candidate pair $(S_i, T_{\hat{\jmath}(i)})$. TRACE then tests whether similarity improves when the match is expanded to a contiguous target span, which captures sentence splitting:

$$\text{sim}\left(S_i, T_{j-\ell} \oplus \cdots \oplus T_j \oplus \cdots \oplus T_{j+r}\right),$$

where $\oplus$ denotes sentence concatenation, and ℓ and r are nonnegative integers constrained by target-sentence boundaries:

$$1 \leq j - \ell \leq j \leq j + r \leq N.$$

TRACE also tests whether similarity improves when a contiguous source span is aligned to a single target sentence, which captures sentence merging:

$$\text{sim}\left(S_{i-\ell} \oplus \cdots \oplus S_i \oplus \cdots \oplus S_{i+r}, T_j\right),$$

where ℓ and r are constrained by source-sentence boundaries:

$$1 \leq i - \ell \leq i \leq i + r \leq M.$$

Adjacent sentences are added only when similarity improves, keeping the expansion local and conservative. Candidate matches are retained only when their similarity exceeds a predefined threshold θ. In the current implementation, $\theta = .60$.[1] Conflicts are resolved deterministically by retaining the highest-scoring candidate and using the leftmost candidate to break ties. The output of Step 1 is a set of matched source-target spans and unmatched source or target spans.

3.3 Step 2: Span-Level Operation Inference

Step 2 converts the matched and unmatched spans from Step 1 into an ordered sequence of revision operations. Matched spans serve as structural anchors. If a matched span preserves its relative order across the two drafts, it is labeled *preserve*. If its relative order changes, it is labeled *move*. Source-only spans are labeled *delete*, and target-only spans are labeled *insert*.

TRACE processes matched spans in their source-text order and emits source-only and target-only spans around these anchors. When multiple structural interpretations are possible, TRACE favors the configuration that minimizes the number of *move* operations. This conservative rule reduces over-identification of movement in ambiguous cases.

The *replace* operation is inferred after the initial sequence is built. Adjacent *delete* and *insert* operations with no intervening operation are compacted into a single *replace* operation. This post-hoc strategy treats replacement as a local relationship between removed and added content while avoiding false replacement labels in regions shaped by larger movement or restructuring.

[1] The threshold is used as a fixed default based on pilot testing. It is user-adjustable and can be modified for applications that require different levels of alignment sensitivity.

3.4 Step 3: Token-Level Revision Analysis

Step 3 refines aligned spans into token-level edits. This step is applied to *preserve*, *move*, and selected *replace* operations. Pure *insert* and *delete* operations are carried forward as span-level rows because they already represent unmatched text.

For an aligned source-target span, let

$$A = (a_1, \ldots, a_m), \qquad B = (b_1, \ldots, b_n)$$

denote the corresponding token sequences. TRACE applies Levenshtein backtrace to identify local *preserve*, *insert*, *delete*, and *replace* edits. Consecutive token edits with the same label are merged into spans, and local token offsets are converted back to global character indices:

$$[p_i^{src}, p_j^{src}]_{\text{global}} = p_0^{src} + [p_i, p_j]_{\text{local}},$$
$$[p_i^{tgt}, p_j^{tgt}]_{\text{global}} = p_0^{tgt} + [p_i, p_j]_{\text{local}}.$$

Move operations are represented in two layers. First, TRACE emits a macro-level *move* row indicating that a source span has been relocated. Second, if the moved span also contains internal changes, TRACE emits token-level rows describing those local edits. This separates structural reordering from within-span revision.

For *replace* operations, token-level decomposition is applied only when the source and target spans share sufficient surface similarity:

$$\text{sim}(A, B) \geq \gamma.$$

In the current implementation, $\gamma = .40$.[2] If similarity falls below this threshold, TRACE keeps the replacement as a holistic span-level rewriting rather than forcing an arbitrary token alignment.

The final output is an ordered, mixed-granularity representation that captures both global restructuring and local textual edits. Each row includes an operation label, source and/or target text, and global character indices, enabling human inspection and revision analysis. Appendix A provides an illustrative example showing how TRACE compares a source draft and revised target text to produce a mixed-granularity representation of revision operations.

The demonstration web app for TRACE can be accessed here: https://trace.lei-tech.org.

[2]The value $\gamma = .40$ was set based on pilot testing and is user-adjustable.

4 Evaluation Experiment and Procedure

To evaluate TRACE, we compared its revision representations with process-level evidence captured through keystroke logging. This validation examined whether revisions inferred from source-target draft comparison corresponded to revision actions that occurred during writing and produced observable changes in the final revised text.

The study used data from 25 participants, including 16 females and 9 males. Participants revised essays in a custom web-based writing application with integrated keystroke logging. Each participant revised five essays, resulting in 125 source-target draft pairs. The essays ranged from 200 to 500 words and were sampled from the KLiCKe corpus (Tian et al., 2025). Participants were instructed to revise each essay in ways that reflected their natural writing and revision practices. They were asked to make at least five revisions per essay and could make any type of textual change, including addition, deletion, substitution, and reordering.

The writing application adapted FlexKeyLogger (Tian and Crossley, 2026), which identifies five operation types: input, delete, paste, replace, and text move. The raw logs were first cleaned to remove actions that did not directly change the text, such as mouse clicks and function-key presses. The cleaned logs were then filtered to retain only revision operations that produced observable changes in the final draft. This step was necessary because keystroke logs may include transient actions, such as text that was typed and later deleted, or intermediate edits absent from the final product. The filtered logs were then reconstructed by combining adjacent character-level insertions or deletions into larger revision chunks to improve readability and facilitate comparison with TRACE outputs.

Two trained experts independently compared TRACE-identified revisions with revision operations reconstructed from the processed FlexKeyLogger logs. The comparison was conducted bidirectionally for each draft pair. First, the experts examined whether each TRACE-identified revision was supported by corresponding evidence in the logs. Second, they examined whether each final-product revision retained in the logs was captured in the TRACE output. These two directions addressed whether TRACE introduced unsupported revision operations or missed revisions observable in the process data.

Each revision was assigned one of three labels.

A revision was labeled *aligned* when it had a corresponding match in the comparison data. Because TRACE and keystroke logs may segment the same revision differently, alignment was judged primarily by span correspondence rather than exact operation-label identity. Related *delete* and *insert* operations were treated as equivalent to a single *replace* operation when they represented the same textual change. Similarly, a TRACE *move* operation was treated as aligned when the logs showed the same text span being cut or deleted from one location and inserted or pasted elsewhere. A revision was labeled *partially aligned* when a corresponding revision was present but only partially matched. A revision was labeled *not aligned* when no corresponding revision could be identified. The two experts reached complete agreement on the alignment labels for both comparison directions.

5 Results

Table 1 presents the number and types of revision operations identified by TRACE and reconstructed from the filtered keystroke logs. TRACE identified 2,022 revision operations, whereas the processed logs contained 2,081 revision operations. The distributions differed across operation types because TRACE and FlexKeyLogger represent revisions at different levels of granularity. For example, TRACE may compact related deletion and insertion spans into a single *replace* operation, whereas keystroke logs may record them as separate process events. In addition, input and paste operations in FlexKeyLogger were combined as *insert* operations to align with the TRACE label set.

One notable difference concerns *move* operations. TRACE identified 11 *move* operations, whereas no operations were recorded as *move* in the processed keystroke logs. This difference reflects how FlexKeyLogger defines and detects movement. In FlexKeyLogger, *move* refers to a mouse-based operation in which a chunk of text is selected and dragged to another location. In contrast, the movements identified by TRACE in this dataset were carried out through keyboard-mediated actions, such as cutting or deleting text from one location and then pasting or inserting it elsewhere. These actions were therefore represented in the keystroke logs as deletion, paste, or insertion events rather than as *move* operations.

Table 2 presents the results of the bidirectional alignment between TRACE operations and the processed keystroke-log evidence for the 125 draft pairs. In the first row, TRACE operations were labeled using keystroke logs as the reference; in the second row, keystroke-log operations were labeled using TRACE operations as the reference. All TRACE-reported operations were verified in the processed logs, and all final-product revisions reconstructed from the logs were represented in TRACE outputs. This finding provides initial evidence that TRACE can accurately recover product-relevant revision operations from source-target draft comparison.

Revision Operation	TRACE	Keystroke Logs
Delete	454	655
Insert	510	572
Replace	1,047	854
Move	11	0
Total	2,022	2,081

Table 1: Revision operations identified by TRACE and reconstructed from keystroke logs.

Operation Set	Reference	Aligned	Partial	None
TRACE operations	Keystroke logs	100%	0%	0%
Keystroke logs	TRACE operations	100%	0%	0%

Table 2: Bidirectional alignment between TRACE operations and processed keystroke-log operations. Each row indicates which operation set was labeled and which source was used as the reference for comparison. *Partial* indicates partially aligned operations, and *None* indicates operations with no corresponding match in the reference source.

6 Conclusion

This paper introduced TRACE, an automated framework for estimating interpretable revision operations from source-target draft pairs. TRACE addresses a methodological need in revision analysis by representing textual change across multiple levels of granularity. Rather than producing only low-level edit scripts or sentence-level correspondences, TRACE integrates span-level matching, operation inference, and token-level refinement to identify preserved text, insertions, deletions, replacements, and moved spans while maintaining links to the original source and target texts.

The evaluation compared TRACE outputs with processed keystroke-log evidence from 125 revised drafts. The results showed complete bidirectional alignment between TRACE-identified operations and final-product revisions reconstructed from the

logs. This finding provides initial evidence that revision operations inferred retrospectively from draft pairs can correspond closely to observable writing-process actions, at least for revisions that remain visible in the final revised text.

TRACE contributes to writing analytics, revision research, and AI-mediated writing studies by offering a reusable representation of textual change that does not require specialized process data. In contexts where writers revise their own drafts, adapt AI-generated text, or respond to AI writing feedback, TRACE can support systematic analysis of what changed and where those changes occurred. More broadly, the framework provides a methodological foundation for future work on revision purpose, revision quality, and human-AI collaboration in writing.

7 Limitations

This study has two main limitations. First, the evaluation validates TRACE against final-product revision evidence reconstructed from keystroke logs, rather than against the full composing process. The processed keystroke-log data retained only revision actions that produced observable changes in the final revised draft. As a result, revisions that were made during writing but later undone, overwritten, or removed were excluded from the comparison. As TRACE relies on the relationship between the initial and final drafts, it is designed to estimate revisions that remain visible in the final text, not transient process actions that leave no observable trace.

Second, although TRACE provides an interpretable multi-granularity representation of revision operations, the operation sequence does not always correspond directly to higher-level revision units or edit intents. A single intended revision may be represented through multiple related operations. Therefore, additional grouping or aggregation may be needed when TRACE outputs are used for revision classification or analysis of revision quality. Future work can build on TRACE by developing methods for clustering related operations into larger revision units that better correspond to writers' rhetorical or communicative intentions.

8 Acknowledgments

The research reported here was supported by the Institute of Education Sciences, U.S. Department of Education, through Grants R305A180261, R305N210041, and R305T240035 to Arizona State University. The opinions expressed are those of the authors and do not represent views of the Institute or the U.S. Department of Education.

References

Regina Barzilay and Noemie Elhadad. 2003. Sentence alignment for monolingual comparable corpora. In *Proceedings of the 2003 conference on Empirical methods in natural language processing*, pages 25–32.

Amit Bronner and Christof Monz. 2012. User edits classification using document revision histories. In *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics*, pages 356–366.

Behice Ceyda Cengiz, Zeynep Bilki, Amine Hatun Ataş, and Berkan Celik. 2025. Exploring second language writers' engagement with chatgpt feedback: Revision behaviors and perceptions. *System*, 134:103837.

Johannes Daxenberger and Iryna Gurevych. 2013. Automatically classifying edit categories in wikipedia revisions. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 578–589.

Lester Faigley and Stephen Witte. 1981. Analyzing revision. *College Composition & Communication*, 32(4):400–414.

Ying Gao, Christian Dieter D Schunn, and Qiuying Yu. 2019. The alignment of written peer feedback with draft problems and its impact on revision in peer assessment. *Assessment & Evaluation in Higher Education*, 44(2):294–308.

Homa B Hashemi and Christian D Schunn. 2014. A tool for summarizing students' changes across drafts. In *International Conference on Intelligent Tutoring Systems*, pages 679–682. Springer.

Chao Jiang, Wei Xu, and Samuel Stevens. 2022. arxivedits: Understanding the human revision process in scientific writing. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 9420–9435.

Svetlana Koltovskaia, Payam Rahmati, and Hooman Saeli. 2024. Graduate students' use of chatgpt for academic text revision: Behavioral, cognitive, and affective engagement. *Journal of Second Language Writing*, 65:101130.

VI Lcvenshtcin. 1966. Binary coors capable or 'correcting deletions, insertions, and reversals. In *Soviet physics-doklady*, volume 10, pages 707–710.

Mariëlle Leijten and Luuk Van Waes. 2013. Keystroke logging in writing research: Using inputlog to analyze and visualize writing processes. *Written Communication*, 30(3):358–392.

Hui-Tzu Min. 2006. The effects of trained peer review on efl students' revision types and writing quality. *Journal of second language writing*, 15(2):118–141.

Eugene W Myers. 1986. An o (nd) difference algorithm and its variations. *Algorithmica*, 1(1):251–266.

Yu Tian and Scott Crossley. 2026. Keystroke logging. In *Digital and Internet-Based Research Methods in Applied Linguistics*, pages 310–335. John Benjamins Publishing Company.

Yu Tian, Scott Crossley, and Luuk Van Waes. 2025. The klicke corpus: Keystroke logging in compositions for knowledge evaluation. *Journal of Writing Research*, 17(1):23–60.

Yu Tian and Sara T Cushing. 2025. Exploring the application of keystroke logging techniques to research in second language (l2) writing. *Research Methods in Applied Linguistics*, 4(1):100179.

Robert A Wagner and Michael J Fischer. 1974. The string-to-string correction problem. *Journal of the ACM (JACM)*, 21(1):168–173.

Fan Zhang and Diane Litman. 2014. Sentence-level rewriting detection. In *Proceedings of the Ninth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 149–154.

A TRACE Revision Analysis Example

Below is an illustrative example of how TRACE represents revisions by comparing a source draft with a revised target text. In this example, the writer relocates the sentence about dogs, inserts the word "also" within the moved sentence, deletes part of the sentence about cats, and adds a new final question. The TRACE output represents these changes through a mixed-granularity sequence of revision operations, including span-level movement and token-level insertions, deletions, and preserved text segments.

Source text I like dogs because they are friendly and loyal. I prefer cats because they are quiet and independent.

Target text I prefer cats because they are quiet. I also like dogs because they are friendly and loyal. What is your preference?

ID	Operation	Source Text	Source Index	Target Text	Target Index
1	MOVE	I like dogs because they are friendly and loyal.	[0, 47]	I like dogs because they are friendly and loyal.	[38, 85]
2	PRESERVE	I	[0, 0]	I	[38, 38]
3	INSERT	–	–	also	[40, 43]
4	PRESERVE	like dogs because they are friendly and loyal.	[2, 47]	like dogs because they are friendly and loyal.	[45, 90]
5	PRESERVE	I prefer cats because they are quiet	[49, 84]	I prefer cats because they are quiet	[0, 35]
6	DELETE	and independent	[86, 100]	–	–
7	PRESERVE	.	[101, 101]	.	[36, 36]
8	INSERT	–	–	What is your preference?	[92, 115]

Table 3: Example TRACE output illustrating mixed-granularity revision representation. The macro-level MOVE operation captures the relocation of the sentence about dogs, while subsequent rows identify local token-level changes within aligned spans, including the insertion of "also," the deletion of "and independent," preserved text segments, and the insertion of the final question. Source and target indices indicate inclusive character positions in the original source and target texts.

Comparing θ Representation Strategies for Reconstructing Item Characteristic Curves with LLMs

Sungjin Nam
ACT Education Corp.
500 Act Dr
Iowa City, IA 52243
sungjin.nam@act.org

Abstract

We examine how strategies for representing student skill (θ) affect IRT-3PL item-parameter recovery from reconstructed item characteristic curves. We compare natural-language descriptors with signed decimal and scientific-notation anchors across anchor counts and placements. Our results show that dense numeric grids improve recovery, while non-uniform placements produce parameter-specific trade-offs, showing that θ representation matters.

1 Background

Collecting student responses for item field testing can be costly. To address this issue, previous studies have suggested using predicted item parameters to reduce calibration sample requirements (Ulitzsch et al., 2025), or recovering item parameters directly from neural text features or item metadata (Yancey et al., 2024; Kapoor et al., 2025).

Recent work has instead used generative models to predict response probabilities across student ability levels. Ormerod (2026) reframed item parameter prediction as item characteristic curve (ICC) reconstruction, predicting $p(\text{correct})$ at a set of natural-language θ anchors and deriving item parameters by fitting the IRT-2PL logistic function. Scarlatos et al. (2025) used single-digit numeric θ values to simulate open-ended responses and recover a 1PL parameter. However, these studies leave open how θ should be represented, how many anchors are needed, and whether ICC reconstruction extends to 3PL passage-based MCQs, where recovery depends on slope (a), location (b), and lower-asymptote (c) behavior.

2 Aims

This study extends LLM-based ICC reconstruction to 3PL parameters (a, b, c) on high-stakes, passage-based reading comprehension items, rather than treating θ representation as a fixed design choice. Specifically, we ask how IRT-3PL recovery performance changes with:

- RQ1. θ **encoding types**: Descriptor, Decimal, or P10 notation.
- RQ2. **Anchor counts**: 7 to 35 θ anchors.
- RQ3. **Anchor placements**: Uniform, normal-quantile, or adaptive-local spacing.

3 Samples

The dataset contains 1,720 passage-based reading MCQs from 43 high-stakes operational forms for assessing high school students. Each form included four passages and 40 items. Each passage was about 700 words long and was linked to 10 items. Reference IRT-3PL parameters were calibrated and equated using operational responses. We used 5-fold cross-validation grouped by form.

4 Methods

4.1 Training

We fine-tuned Qwen-3-1.7B with LoRA, using rank 16, $\alpha = 128$, learning rate 5×10^{-5}, and weight decay 0.03. The baseline followed Ormerod (2026) and used 20 natural-language descriptors for θ.

Each input combined the passage, item stem, options, and a θ anchor in $[-4, 4]$. Targets were $p(\text{correct} \mid \theta)$ values from the reference 3PL ICC. The model output was the probability assigned to the correct-answer option token, and the model was trained with cross-entropy loss.

Training batches grouped multiple θ anchors per item and multiple items from the same passage, allowing the model to observe within-item ICC variation and within-passage item structure. Prompt caching reduced duplicate computation over the system prompt, passage, and item text. In this setup, the LLM acts as a curve estimator. It first

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 52–56
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

predicts response probabilities, and item parameters are recovered afterward from the reconstructed ICC.

4.2 θ Encoding types

We compared three θ encoding types. The *Descriptor* baseline used 20 ability labels from Ormerod (2026), ranging from "Critical" to "Exemplary," which are interpretable but difficult to expand while preserving distinct meanings. *Decimal* encoding used signed numeric strings, such as `+2.35`, to leverage the model's prior knowledge of numbers. *P10* notation converted values into special scientific-notation tokens for sign, digits, and exponent, such as $1.39 \rightarrow$ `<+><1><3><9><E-2>`, following Akhauri et al. (2025). This tested whether explicit numeric tokenization offsets LLM weaknesses in handling continuous numbers (Golkar et al., 2024). Unlike descriptors, both numeric formats can add anchors mechanically without inventing new ability labels.

4.3 θ Counts

Following grid-size studies in IRT-3PL recovery (Casabianca and Lewis, 2015), we tested 7, 14, 21, 28, and 35 θ anchors. This range compares the descriptor-like budget of 20 anchors with denser numeric grids.

4.4 θ Placements

The default placement used uniform spacing over $[-4, 4]$ with single-decimal-place θ values.

At the 35-anchor setting, we also tested *normal-quantile* spacing, with 33 anchors in $[-1.89, 1.89]$, and *adaptive-local* spacing, with 14 training anchors near each item's b plus 21 uniform anchors. Held-out adaptive-local inference used the fixed uniform grid. Both non-uniform settings used two-decimal-place θ values. These conditions tested whether allocating more anchors near the high-density or high-slope regions of the ICC improves recovery.

4.5 Parameter Recovery

For each item, predicted $p(\text{correct} \mid \theta)$ values formed a reconstructed ICC. We fit the 3PL curve by least-squares curve fitting:

$$P(\theta) = c + \frac{1 - c}{1 + \exp\left[-1.7a(\theta - b)\right]}. \quad (1)$$

This yielded estimates of (a, b, c), which we compared with the reference parameters using fold-mean Pearson r and RMSE.

5 Results

Table 1 shows that the strongest models used numeric anchors. However, no single θ design dominated every parameter. Advantages depended on anchor count and placement.

5.1 Numeric θ Encodings Are Comparable to Language Descriptors (RQ1)

At comparable anchor counts of 20 and 21, Decimal and P10 were mixed relative to the Descriptor baseline. They were marginally better for a by RMSE and for c by Pearson r and RMSE, but worse for b by Pearson r and RMSE (Table 1). This pattern also appeared at 7 and 14 anchors. Decimal usually outperformed P10, suggesting that using familiar numeric strings was easier than learning new notation tokens in this training setting.

5.2 More θ Anchors Improve Overall Parameter Recovery (RQ2)

The Descriptor baseline was stronger for b at lower anchor counts, from 7 to 21 anchors, but numeric encodings improved as the θ grid became denser. At 28 and 35 anchors, Decimal and P10 improved recovery for nearly all 3PL parameters. Thus, numeric θ anchors may be less efficient per anchor but easier to scale, which matters when ICC fidelity depends on denser curve coverage.

5.3 Nonlinear θ Placements with Trade-offs (RQ3)

Normal-quantile placement improved a and b recovery but weakened c, consistent with concentrating anchors near $\theta = 0$ and leaving sparse tail coverage. Adaptive-local placement improved only a, indicating training-time local anchors did not transfer consistently to uniform-grid inference. The placement results suggest that central anchors help slope and difficulty, while tail coverage remains important for guessing.

6 Conclusions

In this study, we show that numeric θ encodings can match or exceed language descriptors for 3PL parameter recovery when enough anchors are used. Numeric encodings provide a useful middle ground between leveraging tokens already learned by the

pretrained language models and enabling more finely spaced θ grids. Dense numeric grids improved overall recovery, while non-uniform placement produced parameter-specific trade-offs. For this dataset, the best default is a dense Decimal grid, with normal-quantile spacing useful when a and b are prioritized over c. Future work should test larger generative models, reasoning traces, and domains beyond reading comprehension.

References

Yash Akhauri, Bryan Lewandowski, Cheng-Hsi Lin, Adrian N. Reyes, Grant C. Forbes, Arissa Wongpanich, Bangding Yang, Mohamed S. Abdelfattah, Sagi Perel, and Xingyou Song. 2025. Performance prediction for large systems via text-to-text regression. *arXiv preprint arXiv:2506.21718*.

Jodi M. Casabianca and Charles Lewis. 2015. Irt item parameter recovery with marginal maximum likelihood estimation using loglinear smoothing models. *Journal of Educational and Behavioral Statistics*, 40(6):547–578.

Siavash Golkar, Mariel Pettee, Michael Eickenberg, Alberto Bietti, Miles Cranmer, Geraud Krawezik, Francois Lanusse, Michael McCabe, Ruben Ohana, Liam Parker, Bruno Régaldo-Saint Blancard, Tiberiu Tesileanu, Kyunghyun Cho, and Shirley Ho. 2024. xval: A continuous numerical tokenization for scientific language models. *arXiv preprint arXiv:2310.02989*.

Radhika Kapoor, Sang T Truong, Nick Haber, Maria Araceli Ruiz-Primo, and Benjamin W Domingue. 2025. Prediction of item difficulty for reading comprehension items by creation of annotated item repository. *arXiv preprint arXiv:2502.20663*.

Christopher Ormerod. 2026. Reconstructing item characteristic curves using fine-tuned large language models. *arXiv preprint arXiv:2601.02580*.

Alexander Scarlatos, Nigel Fernandez, Christopher Ormerod, Susan Lottridge, and Andrew Lan. 2025. Smart: Simulated students aligned with item response theory for question difficulty prediction. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 25071–25094.

Esther Ulitzsch, Dmitry Belov, Oliver Lüdtke, and Alexander Robitzsch. 2025. Using item parameter predictions for reducing calibration sample requirements—a case study based on a high-stakes admission test. *Journal of Educational Measurement*.

Kevin P. Yancey, Andrew Runge, Geoffrey Laflair, and Phoebe Mulcaire. 2024. Bert-irt: Accelerating item piloting with bert embeddings and explainable irt models. In *Proceedings of the 19th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2024)*, pages 428–438.

A Parameter Recovery and ICC Reconstruction Results

Table 1: Recovered IRT-3PL parameters by θ encoding and anchor count. Values are held-out fold means across 5-fold passage-grouped cross-validation. The best value in each column is shown in bold, and values better than the descriptor baseline are underlined. *Uniform* placement is evenly spaced over $[-4, 4]$; *normal-quantile* places 33 of 35 anchors in $[-1.89, 1.89]$; and *adaptive-local* clusters training anchors around each item's reference b, with fixed uniform anchors used for held-out inference. Average columns are simple averages across a, b, and c.

Encoding	Placement	θ count	$r(a)$	$r(b)$	$r(c)$	Avg. r	RMSE(a)	RMSE(b)	RMSE(c)	Avg. RMSE
Decimal	*uniform*	7	0.548	0.461	0.451	0.487	0.289	0.882	0.071	0.414
		14	0.531	0.461	0.457	0.483	0.270	0.881	0.084	0.412
		21	0.499	0.500	0.489	0.496	0.288	0.825	**0.065**	0.393
		28	0.542	0.515	0.489	0.515	0.280	0.831	0.068	0.393
		35	0.547	0.523	**0.499**	**0.523**	0.270	0.812	0.068	0.383
	normal-quantile	35	0.551	**0.535**	0.477	0.521	**0.265**	**0.804**	0.075	**0.381**
	adaptive-local	35	**0.562**	0.498	0.382	0.481	0.521	1.113	0.129	0.588
P10	*uniform*	7	0.533	0.459	0.429	0.474	0.286	0.911	0.079	0.425
		14	0.523	0.455	0.450	0.476	0.281	0.884	0.080	0.415
		21	0.507	0.513	0.469	0.496	0.292	0.822	0.067	0.394
		28	0.533	0.523	0.469	0.508	0.288	0.819	0.074	0.394
		35	0.543	0.527	0.470	0.513	0.272	0.816	0.067	0.385
	normal-quantile	35	0.504	0.516	0.479	0.500	0.279	0.819	0.068	0.389
	adaptive-local	35	0.542	0.478	0.402	0.474	0.503	0.953	0.121	0.526
Descriptors (Ormerod, 2026)	*uniform*	20	0.541	0.513	0.443	0.499	0.294	0.820	0.073	0.396

Table 2: ICC reconstruction by θ encoding and anchor count. Held-out fold means across 5-fold passage-grouped cross-validation. Common-grid metrics interpolate each predicted ICC to $\theta \in [-3.75, 3.75]$ at 0.05 increments before computing item-level errors. Lower is better. Best in each column is shown in bold based on unrounded values; values better than the descriptor baseline are underlined. Rows use the same model settings as Table 2.

Encoding	Placement	θ count	ICC RMSE	ICC MAE	Normal-wtd. RMSE	Observed-grid RMSE
Decimal	*uniform*	7	0.101	0.084	0.129	0.094
		14	0.102	0.084	0.127	0.099
		21	0.096	0.077	0.121	0.093
		28	0.094	0.077	0.120	0.092
		35	**0.092**	**0.075**	0.116	**0.089**
	normal-quantile	35	0.093	0.076	**0.116**	0.114
	adaptive-local	35	0.145	0.130	0.143	0.141
P10	*uniform*	7	0.104	0.086	0.132	0.097
		14	0.102	0.084	0.128	0.099
		21	0.095	0.077	0.120	0.092
		28	0.093	0.076	0.118	0.092
		35	0.094	0.077	0.120	0.092
	normal-quantile	35	0.093	0.077	0.117	0.116
	adaptive-local	35	0.148	0.132	0.144	0.146
Descriptors (Ormerod, 2026)	*uniform*	20	0.094	0.078	0.118	0.093

Assessing Large Language Model Performance in Post-Certification-Examination Comment Categorization

Huaping Sun, Ph.D., MBA
Kristin O'Brien, M.Ed.
Colleen B. Kave, J.D.
David Shin, Ph.D.
Qiao Lin, Ph.D.
Jeffrey M. Lyness, M.D.
American Board of Psychiatry and Neurology
Correspondence: hsun@abpn.org

Abstract

This study evaluates a large language model (LLM) for automating the analysis of post-certification-examination comments. A total of 1,406 Neurocritical Care certification examination comments were independently coded by two human raters for sentiment (*positive*, *negative*, and *mixed*) and nine predefined thematic categories. An LLM performed the same tasks under three conditions: baseline prompting, thematic definitions and coding rules, and thematic definitions plus human-coded training examples. Human raters demonstrated high agreement across sentiment and thematic categories. Baseline LLM-human agreement was generally moderate. Contrary to expectations, providing thematic definitions and coding rules decreased LLM-human agreement. Human-coded examples substantially improved LLM performance, especially in concrete, well-defined domains. These findings suggest that LLMs may support preliminary coding of examination feedback, but human review remains essential for complex and subjective comments.

1 Introduction

The American Board of Psychiatry and Neurology (ABPN, Deerfield, IL) provides candidates the opportunity to submit free-response comments immediately after completing certification examinations. These comments offer valuable feedback regarding examination content, administration, and overall testing experience. Over time, such feedback has informed content revisions and supported ongoing efforts to improve examination quality.

Despite their value, systematic analysis of post-examination comments presents several operational challenges. Traditionally, ABPN staff manually review and code comments for sentiment and thematic categories (e.g., question quality, content balance, testing platform functionality). This process is time-consuming, labor-intensive, and subject to variability in human judgment. As the volume of candidate feedback continues to increase, there is a growing need for more scalable, efficient, and consistent analytical approaches to comment analysis.

Recent advances in artificial intelligence, particularly large language models (LLMs), offer a promising opportunity to automate aspects of qualitative feedback analysis (Bennis & Mouwafaq, 2025; Parfenova et al., 2024). This study adopts a progressive evaluation framework to examine the feasibility and accuracy of using an LLM to classify post-examination comments by sentiment and thematic category. Specifically, the study investigates the extent to which LLM-generated classifications align with human judgments and whether additional guidance or training examples improve LLM coding performance.

2 Methods

2.1 Data Sample

The study sample included 1,406 free-text comments submitted through an optional post-examination survey following three Neurocritical Care (NCC) certification examination administrations: 587 comments from 2021, 350 from 2022, and 469 from 2024.

2.2 Coding Framework

Each comment was coded with respect to two dimensions: sentiment and thematic category. Sentiment analysis classified each comment as *positive*, *negative*, or *mixed*. Thematic categorization involved identifying the presence or absence of one or more of nine predefined categories. These categories were developed

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 57–61
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

through a preliminary review of NCC post-examination comments and were selected based on their relevance to future examination improvement (Table 1).

Two human raters independently reviewed an initial subset of comments to familiarize themselves with the coding framework and then met to discuss and align their interpretation of the coding criteria (Saldaña, 2025). Following this calibration process, each rater independently coded all comments from the three examination years. The independent ratings were used to assess inter-rater reliability and to establish a benchmark for evaluating LLM performance.

After completing the initial coding, the same two raters re-examined the 2021 comments and resolved discrepancies through discussion. The resulting consensus ratings served as the training dataset used in the final stage of LLM evaluation.

Category	Description
Prefer Time by Section	Desire for the exam to be timed by individual sections rather than pooled time across all sections.
Not Enough Break Time	Insufficient break periods; delays due to longer-than-expected check-in and check-out processes.
Not Enough Total Time	Feeling rushed or lacking adequate time to complete the examination.
Poor Media Quality	Issues with images or videos, such as poor resolution, darkness, lack of clarity, or confusing visuals.
Editorial Errors	Complaints about misspellings, grammatical errors, or duplicate answer choices (distractors).
Poor Question Quality	Questions perceived as unclear, overly vague or detailed, having multiple plausible answers, or lacking well-designed answer choices.
Lack of Content Balance	Feedback about disproportionate coverage of certain topics, either too much or too little.
Questions Not Reflecting Practice	Concerns that questions do not mirror real-world clinical practice.
Exam Functions Need Improvement	Suggestions for better usability features, such as easier ways to highlight, mark, or retrieve questions for later review.

Table 1. Categories for Neurocritical Care Post-Examination Comments

2.3 LLM Evaluation Conditions

To examine the extent to which increasing levels of guidance improve agreement between LLM-generated classifications and human ratings, the LLM was evaluated under three conditions:

1) Baseline condition. The LLM performed sentiment analysis and thematic categorization using a minimal prompt, with no additional information regarding category definitions or coding rules.
2) Thematic-definition-and-coding-rule condition. In addition to the base prompt, the LLM was provided with descriptions of each thematic category and coding rules linking sentiment classifications to thematic categories. Specifically, because all thematic categories represented concerns, criticisms, or areas for improvement, category tags were applied only to negative comments or to the negative component of mixed comments.
3) Thematic-input-plus-training condition. In addition to thematic definitions and coding rules, the LLM was provided with the consensus ratings of the 2021 comments for training purposes. This condition was designed to evaluate whether exposure to domain-specific, human-rated examples would further improve LLM coding performance.

We hypothesized that coding accuracy would improve following the introduction of thematic guidance and would improve further after the inclusion of human-coded training examples.

2.4 LLM Implementation

Microsoft Copilot was used as the platform for implementing the LLM-based coding workflow. For each examination year, an Excel file containing one comment per row and one category per column was submitted to the LLM, which generated a new file that preserved the original data structure while adding sentiment classifications and thematic category indicators (1 = present, 0 = absent). Coding was performed separately under each of the three evaluation conditions.

2.5 Statistical Analysis

Agreement was assessed between the two human raters and between the LLM and each human rater under each evaluation condition. Cohen's kappa coefficient (κ), which accounts for agreement expected by chance, was calculated for sentiment classification and for each thematic category separately by examination year. Kappa values were then averaged across examination years and other relevant analytical groups (e.g., raters or categories), as appropriate. Kappa values were interpreted according to the criteria proposed by Landis and Koch (1977), with values of 0.21–0.40, 0.41–0.60, 0.61–0.80, and 0.81–1.00 indicating fair, moderate, substantial, and almost perfect agreement, respectively.

3 Results

Inter-rater agreement between the two human raters was consistently high across sentiment analysis and thematic categorization, with average Cohen's κ ranging from 0.68 to 0.97 across categories (Table 2). Human-human agreement was highest for sentiment analysis (average κ = 0.95) and for categories with relatively concrete coding criteria, such as *Not Enough Break Time* (average κ = 0.97), *Prefer Time by Section* (average κ = 0.94), and *Poor Media Quality* (average κ = 0.94). Lower, although still substantial, agreement was observed for more subjective categories, including *Questions Not Reflecting Practice* (average κ = 0.73) and *Exam Functions Need Improvement* (average κ = 0.68).

Category	Human-Human κ	LLM Baseline κ	LLM Thematic Input κ	LLM Training κ‡
Sentiment	.95	.49	.46	.70
Prefer Time by Section	.94	.66	.56	.83
Not Enough Break Time	.97	.70	.51	.78
Not Enough Total Time	.93	.62	.47	.62
Poor Media Quality	.94	.66	.57	.80
Editorial Errors†	.90	.57	.43	.45
Poor Question Quality	.84	.50	.46	.60
Lack of Content Balance	.84	.38	.28	.49
Questions Not Reflecting Practice	.73	.41	.38	.53
Exam Functions Need Improvement	.68	.42	.37	.25
Average across categories	***.87***	***.54***	***.45***	***.61***

Table 2. Average Cohen's kappa coefficients between human raters and the LLM across examination years

* κ = Cohen's kappa coefficient. 0.00–0.20 = slight, 0.21–0.40 = fair, 0.41–0.60 = moderate, 0.61–0.80 = substantial, and 0.81–1.00 = almost perfect agreement.

† Values for Editorial Errors were averaged across 2021 and 2022 because no editorial errors were identified by human raters in 2024.

‡ Only 2022 and 2024 data were used for LLM training calculations because the 2021 coded data were used for training.

Under the baseline condition, agreement between the LLM and human raters was generally moderate (Table 2). For sentiment analysis, the LLM achieved an average κ of 0.49.

Across thematic categories, agreement ranged from average κ = 0.38 (*Lack of Content Balance*) to average κ = 0.70 (*Not Enough Break Time*).

Providing the LLM with thematic definitions and coding rules did not improve performance (Table 2). Instead, agreement decreased across all outcomes, with the largest decreases observed for *Not Enough Break Time* (average κ = 0.70 to 0.51), *Not Enough Total Time* (average κ = 0.62 to 0.47), and *Editorial Errors* (average κ = 0.57 to 0.43). The LLM classified a higher percentage of comments as *positive* than did the human raters under the thematic input condition (Table 3). Because comments classified as *positive* were not assigned thematic categories, this tendency may have contributed to lower thematic-category tagging rates relative to the baseline condition (Table 4).

Rater	Positive (%)	Negative (%)	Mixed (%)
Human	21.0	48.0	30.8
LLM Baseline	45.3	30.7	24.3
LLM Thematic Input	48.0	29.0	23.0
LLM Training*	16.5	50.5	33.0

Table 3. Average percentage of comments classified as *positive*, *negative*, and *mixed* by human raters and the LLM across examination years

* Only 2022 and 2024 data were used because the 2021 consensus data were used for training.

In contrast, the addition of human-coded training examples substantially improved LLM performance for most outcomes (Table 2). Agreement for sentiment analysis increased from average κ = 0.49 under baseline condition to average κ = 0.70 after training. The highest post-training agreement was observed for *Prefer Time by Section* (average κ = 0.83), *Poor Media Quality* (average κ = 0.80), and *Not Enough Break Time* (average κ = 0.78). Improvements were also observed for *Questions Not Reflecting Practice* (average κ = 0.41 to 0.53), *Lack of Content Balance* (average κ = 0.38 to 0.49), and *Poor Question Quality* (average κ = 0.50 to 0.60).

Thematic Category	Human (%)	LLM Baseline (%)	LLM Thematic Input (%)	LLM Training (%)*
Prefer Time by Section	8.7	4.3	3.7	10.0
Not Enough Break Time	4.5	2.7	1.7	5.0
Not Enough Total Time	15.5	11.0	6.3	20.5
Poor Media Quality	23.0	15.0	14.3	23.5
Editorial Errors	1.5	3.3	2.3	3.0
Poor Question Quality	28.5	18.7	17.0	42.5
Lack of Content Balance	13.5	7.3	6.3	24.5
Questions Not Reflecting Practice	18.8	7.7	6.3	31.0
Exam Functions Need Improvement	5.0	5.0	2.7	17.0

Table 4. Average percentage of thematic categories tagged by human raters and LLMs across examination years

* Only 2022 and 2024 data were used because the 2021 consensus data were used for training.

4 Discussion

This study examined the extent to which an LLM could replicate human coding of post-examination comments and whether additional guidance or training examples improved LLM coding performance. Three key findings emerged. First, human raters demonstrated consistently high agreement across sentiment analysis and thematic categorization, supporting the reliability of the coding framework. Second, the LLM achieved moderate agreement under the baseline prompting condition, with performance varying substantially across categories. Third, providing category definitions and coding rules that introduced task dependency did not improve performance, whereas the addition of human-coded examples resulted in substantial gains for sentiment analysis and most thematic categories.

Both human raters and the LLM exhibited higher agreement for concrete and observable categories such as *Not Enough Break Time*, *Prefer Time by Section*, and *Poor Media Quality,* and lower agreement for categories requiring subjective interpretation, including *Questions not Reflecting Practice* and *Exam Functions Need Improvement*. Consistent with the findings of Liu et al. (2025) using GPT-4, these parallel difficulty patterns observed for human raters and the LLM suggest that lower agreement in certain categories may reflect the complexity of the coding task itself rather than model performance alone.

A particularly noteworthy finding was that providing thematic definitions and coding rules actually reduced agreement in most thematic categories. One potential explanation is that the thematic-input condition introduced task dependency between sentiment and category assignments (i.e., not assigning thematic-category tags to comments classified as *positive*), which unintentionally amplified LLM's potential optimism bias. In contrast, the addition of human-coded examples substantially improved agreement across most outcomes, especially in more clearly defined domains.

Several limitations should be noted. The study used comments from one subspecialty examination, a predefined set of thematic categories, and only one set of thematic input and coding rules, and a single training dataset. Other contexts or iterative prompt-development approaches may yield different results. In particular, thematic category definitions and coding rules should have been separate prompts.

In conclusion, this study suggests that LLMs may be useful for preliminary categorization and prioritization of examination comments for further review. However, human oversight remains important for comments involving ambiguity, complexity, contextual reasoning, and subjective judgment. Future studies may investigate the reasoning that LLM provides after coding execution and explore creative approaches to improve LLM coding performance.

References

Issam Bennis and Safwane Mouwafaq. 2025. Advancing AI-driven thematic analysis in qualitative research: A comparative study of nine generative models on cutaneous leishmaniasis data. *BMC Medical Informatics and Decision Making, 25,* 124. https://link.springer.com/content/pdf/10.1186/s12911-025-02961-5.pdf

J. Richard Landis and Gary G. Koch. 1977. The measurement of observer agreement for categorical data. *Biometrics, 33,* 159–174. https://www.jstor.org/stable/2529310

Xiner Liu, Andres Felipe Zambrano, Ryan S. Baker, Amanda Barany, Jaclyn Ocumpaugh, Jiayi Zhang, Maciej Pankiewicz, Nidhi Nasiar, and Zhanlan Wei. 2025. Qualitative coding with GPT-4: Where it works better. *Journal of Learning Analytics, 12*(1), 169–185. https://files.eric.ed.gov/fulltext/EJ1465623.pdf

Angelina Parfenova, Alexander Denzler, and Jörgen Pfeffer. 2024. Automating qualitative data analysis with large language models. *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Student Research Workshop),* 83–91. https://aclanthology.org/2024.acl-srw.17/

Johnny Saldaña. 2025. *The Coding Manual for Qualitative Researchers* (5th ed.). Sage Publications.

Rethinking Pilot Data: Evaluating LLM Synthetic Data for Scale Development

Margarita Olivera-Aguilar, Samuel H. Rikoon, Paul D. Bailey, Michael B. Kruse, Kamal Middlebrook
American Institutes for Research

Abstract

This study evaluated whether LLM-synthetic data can support K-12 instrument development. We generated LLM-synthetic datasets under various prompt conditions for two surveys and one assessment. Our findings indicate that while some prompt conditions successfully reproduced the overall latent structure of the instruments, recovery of item-level parameters was generally inadequate.

1 Introduction

Developing K–12 assessments and surveys meeting rigorous standards (AERA, APA, and NCME, 2014) is resource intensive but essential for understanding what students know and can do. Rapid advances in AI have led researchers to use synthetic data (AI-produced datasets resembling real data) in educational and behavioral studies (e.g., Benedetto et al., 2024; Bisbee et al., 2024; Motoki et al., 2023). In educational assessment, Liu et al. (2025) and Säuberli et al. (2025) reported moderate to strong correlations between item parameters estimated from LLM responses and those estimated from human responses in some academic subjects, but they also showed important limitations, including restricted variability in respondent proficiency. Synthetic data could help streamline instrument development and reduce costs, but important questions remain. Can LLM synthetic data reproduce the psychometric properties observed in human responses? Do synthetic datasets introduce systematic bias across student groups? How can synthetic data be used responsibly as part of validity evidence? These questions are relevant for both self-report surveys and academic assessments. In survey research, LLMs are often asked to generate responses that reflect people's social, emotional, or behavioral characteristics (which might differ from what is represented predominantly in LLM training data). In academic assessments, the challenge is different: LLMs must simulate how students with varying levels of knowledge and different test-taking strategies respond to content-based (i.e., without simply relying on the LLM's vast knowledge base to select the correct answer invariably). Academic assessments also provide an opportunity to examine whether LLMs can reproduce common testing patterns, such as differences in item difficulty, student misconceptions, guessing, and careless responding.

Our primary aim is to evaluate the viability of LLM-generated synthetic data as a tool for supporting early-stage psychometric analyses. Our intention is not to replace human validation samples, but rather to evaluate the use of synthetic data as a tool for refining measures before large-scale data collection. To do so, we tested multiple prompts using ChatGPT 4.0 with the following goals:

1. Determine whether synthetic datasets produce realistic item response distributions,
2. Examine whether synthetic data produces confirmatory factor analysis (CFA) model fit consistent with human data,
3. Compare synthetic and human datasets with respect to patterns of localized CFA misfit, and
4. Evaluate correspondence between item parameter estimates in human vs. synthetic data.

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 62–68
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

A second methodological goal is to develop a more systematic framework for evaluating synthetic data. Drawing on Monte Carlo simulation methods, we generate and analyze multiple synthetic datasets under each prompting condition rather than relying on a single dataset. This approach allows us to evaluate the stability, consistency, and reproducibility of findings across repeated datasets.

2 Sample

We report findings from K–12 student samples collected using two surveys and one academic subject assessments: (1) the BESSI self-management scale with five response options (Soto et al., 2022; N=478 middle school students), (2) a leadership scale ("LEADER"; N=750 intermediate/middle school students) with four response options, and (3) a science assessment ("SCIENCE", N=750 4th grade students) with two response options (correct/incorrect). We have shielded the names of the latter two measures to comply with contractual obligations.

3 Method

3.1 Analysis of human data

We used the lavaan R package (Rosseel, 2012) to fit one-factor CFA models to the human data, treating BESSI items as continuous and LEADER and SCIENCE items as categorical. We constructed two versions for each instrument: (1) "well-fitting" versions (i.e., BESSI-9, LEADER-6, and SCIENCE-12, with 9, 6 and 12 items, respectively) and (2) "poorly fitting" versions (i.e., BESSI-13, LEADER-9, and SCIENCE-16; 13, 9 and 16 items each, with items added from other parts of each larger instrument. Unlike BESSI, whose factor structure and item properties have been established across multiple studies, LEADER and SCIENCE reflect the ambiguity typical of scale development, where the underlying structure is less certain and some items may perform differently across samples. For example, LEADER-6 was formed by items from three different subscales that nevertheless resulted in adequate fit to a unidimensional model; and although the 12-item SCIENCE form showed adequate model fit, several items had loadings at or below .20 in the human sample. This design allowed us to examine whether synthetic data could reproduce psychometric patterns under both well-established and more uncertain measurement conditions.

3.2 LLM synthetic data generation

Using ChatGPT 4.0, we generated synthetic datasets matching the human sample sizes for each instrument. For each respondent, the model received demographic information and survey items and was instructed to generate plausible responses. For the survey instruments, we tested 29 prompting conditions that varied the information provided to the model, including demographic characteristics, respondent personas, trait-level descriptions, response-style tendencies, item-level response probabilities, and instructions about response consistency.

For SCIENCE, we developed 33 prompts based on the set of BESSI and LEADER survey prompts but adapted specifically for academic assessments. Because responses to multiple-choice test items depend on factors such as content knowledge, item difficulty, misconceptions, and guessing, SCIENCE prompts incorporated information such as demographic profiles, expected performance levels, item difficulty, common misconceptions, test-taking effort, attention, and response strategies. Some conditions also included answer keys or expected percent-correct information. When answer keys were provided, prompts explicitly instructed ChatGPT not to simply select the correct answer, but instead to respond as a student with the specified characteristics would be likely to respond

We initially generated 20 datasets per condition and retained those that most closely reproduced the response distributions observed in the human data. Following a Monte-Carlo simulation approach, for the selected conditions, we then generated 500 datasets per condition to reduce the likelihood that findings are driven by random variation or idiosyncratic characteristics of single or small groups of synthetic dataset(s).

3.3 LLM synthetic data analysis

Consistent with common survey and assessment development practices, our analytic procedure focused on evaluating dimensionality and identifying opportunities for scale refinement. We used the lavaan R package (Rosseel, 2012) to fit a one-factor CFA model to each synthetic dataset to assess the intended factor structure. We then fit a

second model removing items with loadings below .20. We used the same models specifications in the human and LLM synthetic data (e.g., model identification constraints, estimator) to ensure the comparability of results.

Analyses were conducted on the final models (i.e., after excluding items with factor loadings below .20). First, we evaluated the fit of a unidimensional CFA model for each dataset using the Comparative Fit Index (CFI), Tucker-Lewis Index (TLI), and Root Mean Square Error of Approximation (RMSEA). We expected acceptable fit for the well-fitting test forms and poorer fit for the extended forms in which additional items were intentionally introduced to induce model misfit. Second, we examined CFA modification indices to determine whether the added misfitting items were consistently identified as sources of local misfit. Finally, we evaluated item parameter recovery by comparing synthetic and human item parameter estimates using four complementary metrics: (a) **absolute difference (AD)**, calculated as the absolute discrepancy between the human and synthetic estimates; (b) **relative difference (RD)**, calculated as the difference between the synthetic and human estimates divided by the human estimate; (c) **mean absolute relative difference (MARD)**, calculated by first taking the absolute value of the relative bias within each replication and then averaging these values across replications; (d) **root mean squared error (RMSE)**, calculated as the square root of the mean squared discrepancy between the human and synthetic estimates across replications and used to quantify overall recovery error while giving greater weight to larger discrepancies; and (e) **Pearson correlations** between the human and synthetic parameter estimates, which assess the extent to which the relative ordering of item parameters is preserved.

4 Results

4.1 Survey results: BESSI and LEADER

Response distributions. In both surveys, only 3 of 29 prompt conditions produced synthetic datasets using all five item response options in BESSI and all four response options in LEADER. The rest of the conditions resulted in synthetic data using only one or two of the item response options. We focused analyses on the datasets within those three conditions: C18 (providing the model with examples of five previous responses, represented by randomly generated numbers), C24 (asking the model to implicitly adopt a position on the latent trait being measured while also providing randomly generated numbers as the previous five responses), and C25 (asking the model to explicitly identify the latent trait underlying the items, specify the simulated person's standing on that trait, and providing random numbers as the previous five responses).

Items with low loadings. In the synthetic BESSI-13 data, one item consistently showed loading values < 0.20 across datasets (i.e., BESSI_136) (Table S1); this item corresponds to a cooperation scale and was one of the four items added to BESSI-9 to create a poor fitting model. In the synthetic LEADER data, by contrast, no single item was universally problematic (three items were dropped in < 3% of datasets).

Model fit. Overall, C24 and C25 resembled model fit patterns observed in human data, with poor model fit in BESSI-13 (Table S2) and LEADER-9 (Table S3), and better fit in BESSI-9 and LEADER-6.

Modification indices (MIs). The results (Table S4) indicate that for BESSI-13, the largest MIs were associated with residual covariances between items that were not part of the original BESSI self-management scale. For BESSI-9, the highest MI involved a pair of conceptually related items (“Think things through carefully” and “Double-check my work”). In the LEADER synthetic data, the largest MIs correspond to items taken from the same subscales.

Item parameter recovery. The results for BESSI (Table S5) indicate that while correlations with human item parameters were moderate to large, AD, RD and MARD had larger than acceptable values (i.e., > .10). Overall, Condition 24 was preferable for loading accuracy, while Condition 18 was preferable for intercept accuracy. Condition 24 produced the most accurate loading estimates for both BESSI forms, with the lowest AD, MARD, and RMSE. Loadings were generally underestimated, especially in Condition 25. Condition 18 produced the most accurate intercept estimates, with the lowest AD,

MARD and RMSE for both forms. Correlations showed a less consistent pattern: Condition 25 best recovered the ordering of BESSI-13 loadings, whereas Condition 18 performed best for BESSI-9.

The results for LEADER (Table S6) showed larger differences and lower (or even negative) correlations than BESSI. LEADER-9 generally showed better item parameter recovery than LEADER-6, particularly under Condition 24 and in terms of Pearson correlations. Intercepts had substantially stronger Pearson correlations than loadings. MARD varied across thresholds, with generally better recovery for t1 thresholds and considerably larger differences for t2 thresholds. Although loadings had lower RMSE, their weak Pearson correlations indicated that relatively low overall error did not correspond to recovery of the human loading pattern.

4.2 Assessment results: SCIENCE

Analysis with 20 replications. Human item percent-correct values ranged from 16.8% to 53.9%, indicating a relatively difficult assessment. In contrast, most prompt conditions produced LLM synthetic datasets with narrower score distributions and higher item percent-correct values, suggesting that the assessment was easier for the AI-generated respondents. Nineteen conditions that exhibited score ranges smaller than 20 percentage points and/or very high item percent-correct values in both forms were excluded from further consideration.

Because generating 500 dataset replications for every condition required substantial computational resources, we used a two-stage approach. First, we evaluated model fit and item parameter recovery using 20 replications for 14 candidate conditions. We then conducted the full 500 dataset generation and analyses for eight conditions (Table 1) that showed the most promising initial results (see the following sections). Results for the remaining conditions will be examined in future work.

Response distributions. Most prompt conditions produced substantially higher item percent-correct values and narrower score distributions relative to the human data (16.8%-53.9% correct), indicating higher proficiency levels among synthetic respondents versus humans (Table S7). Among the conditions examined, Condition 22 most closely approximated the difficulty of the human assessment, whereas the remaining conditions generally yielded considerably higher percent-correct values.

Prompt	Description
C4	Includes the answer key, instructing the LLM to pick familiar-sounding or misleading answers
C5	Includes the answer key, instructing the LLM to pick the number of items to get correct before answering
C21	Includes an ability level for each individual (below average, average, above average)
C22	Includes the answer key, and an ability level for each individual (below average, average, above average)
C26	Same as 22, with values for the theta distribution (mean and standard deviation)
C28	Same as 26, encouraging chain-of-thought reasoning for the LLM to guess between the two most likely correct responses.
C30	Same as 26, with chain-of-thought reasoning around confidence in response based on item difficulty and ability level.
C32	Same as 26, with chain-of-thought reasoning around writing a paragraph explaining the LLM's answering approach.

Table 1. Prompt description for SCIENCE

Items with low loadings. The results (Table S8) show that item q27 (with a negative loading in the human sample) was dropped by over 80% of replications in every condition. Among the evaluated conditions, C5 most closely reproduced the pattern of low-loading items observed in the human data, with a high percentage of replications dropping items with loadings below .20 in the human data. However, C5 also produced loadings below .20 for several items that exceeded .20 in the human sample. In contrast, C4 was more conservative, consistently identifying two items with low loadings in the human data while rarely classifying higher-

loading items as problematic. However, C4 failed to identify several other items with low or negative loadings in the human data.

Model fit. Only C4 replicated the pattern observed in the human data, with SCIENCE_12 showing adequate fit and SCIENCE_16 failing to do so according to TLI with an average value across replications of .886 (Table S9). It should be noted that these fit measures correspond to the models after dropping the items with low loadings, explaining the adequate model fit for SCIENCE_16 in C5.

Modification indices (MIs). Overall, the average modification indices were modest, with only a few item pairs showing mean values above 10. The largest average MIs occurred primarily under Conditions 30 and 32 and involved q20, q21, and q24, where only q24 was an item added to create a poorly fitting model (Table S10). It should be noted that while q24 had a negative loading in the human data, it was rarely dropped for showing loadings below .20 in the LLM synthetic data. Other item pairs generally had lower average MIs and were infrequently flagged, suggesting that evidence of local dependence was limited and concentrated primarily among q20, q21, and q24.

Item parameter recovery. No condition exhibited AD, RD and MARD values under .10 (Table S11). Regarding loadings in SCIENCE_12, C26 produced the best overall parameter recovery, with the lowest AD (0.243), lowest MARD (0.540), and lowest RMSE (0.336), while maintaining a strong correlation with the human item loadings (r = .630). Although C30 achieved the highest correlation (r = .736), it exhibited substantially greater differences and RMSE. For SCIENCE_16, C26 and C32 showed the strongest performance. C32 minimized differences and error, whereas C26 provided a more favorable balance between recovery accuracy and association with the true parameters (r = .461 vs. .359). For loadings, C26 emerged as the most robust overall condition, offering consistently low differences, low RMSE, and comparatively strong parameter recovery.

For threshold recovery (Table S12), performance was generally weak relative to the loading estimates. For SCIENCE_12, C22 minimized AD (0.575), MARD (1.584), and RMSE (0.663). However, the average correlation with human loadings was negative (r = −.350). C28 provided the strongest correlation (r = .311) but had higher differences and RMSE. For SCIENCE_16, C22 again had the lowest AD (0.662) and RMSE (0.722), while C26 had a nearly identical MARD (1.988 versus 1.986 for C22). Nevertheless, both conditions produced negative correlations. C28 yielded the highest positive correlation (r = .078), although this association was very weak.

5 Conclusions

Overall, the findings indicate that developing prompts to yield usable synthetic data is not straightforward, and that AI-generated synthetic respondents are currently more effective for reproducing broad psychometric patterns than for recovering precise item parameters.

In our study, only 3 of the 29 survey-related prompts resulted in usable synthetic data. However, outcomes from these successful cases were encouraging. Under two specific prompting conditions, ChatGPT generated synthetic data useful for supporting early-stage scale development. Synthetic data were generally able to recover broad relationships among items and structural patterns, including distinctions between well-fitting and poorly fitting models. However, reproduction of specific item-level parameters was less favorable.

Performance was notably weaker for the science assessment. Although several conditions reproduced isolated features of the human data (e.g., item percent-correct distributions, identifying items with low loadings in the human data), none produced psychometrically convincing recovery across the full set of evaluation criteria. Across all conditions, threshold recovery was particularly poor, with weak and frequently negative correlations with the human threshold estimates. Thus, while some prompt conditions reproduced selected characteristics of the human data, none generated synthetic datasets that could be considered close

approximations of the human psychometric structure.

Several explanations may account for the relatively weak correspondence between the human and LLM synthetic datasets observed in the present study. First, the results suggest that generating realistic synthetic data may be more challenging for knowledge assessments than for survey scales. The strongest evidence of correspondence between human and synthetic data has been observed for BESSI, a behavioral inventory, whereas the SCIENCE assessment showed substantially poorer recovery of item characteristics, model fit, and response distributions. Similarly, although LEADER did not show the same level of correspondence observed for BESSI, its results were generally more favorable in terms of overall structural patterns than those obtained for SCIENCE. This suggests that AI-generated data may be better able to reproduce responses to attitudinal, behavioral, or self-report survey items than responses to items requiring specific knowledge or competencies. One possible explanation is that survey responses are more directly tied to the semantic and conceptual information available in the training data of large language models, whereas assessment responses require accurate modeling of variations in knowledge, misconceptions, and ability levels across individuals.

At the same time, differences in instrument maturity may also contribute to the observed findings. BESSI is the only instrument examined in this project that has extensive prior psychometric evidence supporting its structure and item functioning. As a result, the item parameters obtained from human respondents may be relatively stable and more likely to generalize across samples. In contrast, LEADER and SCIENCE are newer instruments for which item parameters have been studied less extensively. Consequently, some item characteristics observed in the human sample may be sample dependent rather than stable properties of the instrument. Under this interpretation, at least some of the discrepancies between the human and synthetic datasets may reflect instability in the human parameter estimates rather than failures of the LLM synthetic data. It is possible that the synthetic data are capturing patterns that would emerge in other human samples, even if those patterns do not align perfectly with the specific sample used in the present analysis.

These considerations suggest an important direction for future research. Rather than using newer instruments as the primary benchmark for evaluating prompting strategies, future studies should first examine a broader set of instruments with well-established psychometric characteristics and evidence of parameter stability across multiple samples. Such instruments would provide a stronger basis for identifying the prompting conditions that most effectively reproduce human data. Once effective prompting strategies have been identified using mature instruments, they can then be applied and evaluated on newer instruments such as LEADER and SCIENCE.

Finally, the present study evaluated a particularly stringent use case in which all responses were generated exclusively by the language model. Future work may benefit from hybrid approaches in which AI-generated data augment, rather than replace, human data. For example, a language model could be informed by a smaller human sample and then generate additional respondents consistent with the observed response patterns. Such approaches may prove more effective for reducing data collection requirements while maintaining psychometric properties that more closely resemble those observed in human samples.

References

American Educational Research Association, American Psychological Association, & National Council on Measurement in Education. (2014). *Standards for educational and psychological testing*. https://www.apa.org/science/programs/testing/standards

Benedetto, L., Aradelli, G., Donvito, A., Lucchetti, A., Cappelli, A., & Buttery, P. (2024, November). Using LLMs to simulate students' responses to exam questions. In A-O, Yaser, M. Bansal, & Y-N Chen (Eds.), *Findings of the Association for Computational Linguistics: EMNLP 2024* (pp. 11351–11368). Association for Computational Linguistics.

Bisbee, J., Clinton, J. D., Dorff, C., Kenkel, B., & Larson, J. M. (2024). Synthetic replacements for human survey data? The perils of large language models. *Political Analysis*, *32*(4), 401–416. https://doi.org/10.1017/pan.2024.5

Motoki, F. Y. S., Monteiro, J., Malagueño, R., & Rodrigues, V. (2023). *Large language models in survey research: Generating synthetic data and unlocking new possibilities*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4595896

Liu, Y., Bhandari, S., & Pardos, Z. (2025). Leveraging LLM respondents for item evaluation. A psychometric analysis. *British Journal of Educational Technology, 56*, 1028–1052 https://bera-journals.onlinelibrary.wiley.com/doi/pdf/10.1111/bjet.13570

Säuberli, A., Frassinelli, D., & Plank, B. (2025). Do LLMs give psychometrically plausible responses in educational assessments? https://aclanthology.org/2025.bea-1.21.pdf

Soto, C. J., Napolitano, C. M., Sewell, M. N., Yoon, H. J., & Roberts, B. W. (2022). An integrative framework for conceptualizing and assessing social, emotional, and behavioral skills: The BESSI. *Journal of Personality and Social Psychology, 123*(1), 192–222. https://doi.org/10.1037/pspp0000401

Rosseel, Y. (2012). *lavaan: An R package for structural equation modeling*. *Journal of Statistical Software, 48*(2), 1-36. https://doi.org/10.18637/jss.v048.i02

A Supplementary Material

All tables are included in the Supplementary materials: https://docs.google.com/document/d/1YihzSTljcgox9XF0IUdjDI5236KEARth/edit?usp=sharing&ouid=110493509676671251703&rtpof=true&sd=true

Hierarchical Attention Network Architecture for Automated Response Time Prediction

Wei Schneider
Cambium Assessment

Abstract

This paper presents a novel deep learning architecture for item response time prediction to create human readable feedback for content creators in educational assessment. The model employs a two-tier attention mechanism, mirroring the cognitive hierarchy of reading: the first tier identifies influential words nested within their sentence context, while the second tier evaluates the contribution of individual sentences to the overall item block. By leveraging this specific hierarchical structure, the model improves explainable AI (xAI) in educational measurement AI research. While architecturally simpler models can achieve comparable or superior raw predictive accuracy, the two-tier structure is designed specifically for xAI — to yield transparent, sentence- and world-level attribution that item developers can act on directly. This methodology bridges the gap between opaque, black-box predictive modeling and actionable, human interpretable feedback for instructional design and item development.

1. Introduction

Response time is a critical factor in educational assessment, directly impacting test fairness, validity, and cognitive fatigue (Bridgeman & Cline, 2000; Ackerman & Kanfer, 2009). Furthermore, predicting response time a priori is highly valuable, enabling content writers and test developers to better gauge item limits and pacing required for various test forms.

A growing body of literature has successfully leveraged large language models (LLMs) in prediction of item difficulty and response time (e.g., Baldwin, Yaneva, Mee, Clauser, & Ha, 2021; Veeramani, Thapa, Shankar, & Alwan, 2024), demonstrating the efficacy of natural language processing techniques and pre-trained language models in psychometric modelling. However, many of these approaches often overlook the limitations inherent in these abstract models, which frequently operate as 'black boxes'.

One of the earlier deep learning models that was inherently a good candidate for xAI (Love et al., 2023) was the Hierarchical Attention Network (HAN; Yang et al., 2016) architecture. HAN introduces a two-tiered attention mechanism that models at both the word and sentence levels. By translating complex model weights into human-readable feedback, HAN explains its predictions by highlighting the specific words and sentences that drive its classification decisions. Such attention framework largely increases the transparency of the model but also could give practical feedback on how text features influence predicted feedback. This could be a very good use for item development as a feedback could pinpoint exactly where in an item text specific phrasing correlates with longer student response time.

The primary goal of this study is to leverage a hierarchical attention framework to construct an explainable deep learning model for educational assessment. Specifically, this research pursues a dual purpose: first, to assess the predictive accuracy of a two-tier attention architecture in forecasting student response times from raw item text; and second, to investigate the model's interpretability by examining whether its dual-level attention mechanisms yield actionable, human-readable feedback on textual complexity for item developers.

2. Methods

A sample of 656 middle school science items were selected for this study, comprising 342

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 69–78
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

standalone items and 314 cluster items. The comprehensive item texts--including both the stimulus and questions--was extracted from the assessment item bank through an Application Programming Interface (API). Due to API extraction limitations, embedded figures and tables were either incomplete or missing from the raw data dump. To maintain dataset integrity, systematic text preprocessing was applied to remove incomplete sentence fragments and partial tabular artifacts. To account for the omitted visual and structural components in the items that could affect response time, five proxy variables "number of figures", "number of tables", "number of equations", "number of score points", and "item type" were add to the dataset. The operational target variable for prediction was the 80th percentile of student response time (T_{80}), which was historically established as an administrative benchmark to distinguish regular length items from extra-long items. T_{80} serves as a metric for modeling item-level time intensity for this experimental study. The time distribution is shown in Figure 1. The distribution of response time is a right-skewed distribution where 5% of the entire distribution falls above 13.62 minutes. Notably, the distribution is bimodal, with distinct concentrations around 2-3 minutes and 8-10 minutes. This bimodality reflects the distinction between standalone and cluster items, Standalone items have a median response time of 2.70 minutes, while the cluster items have a median response time of 9.30 minutes.

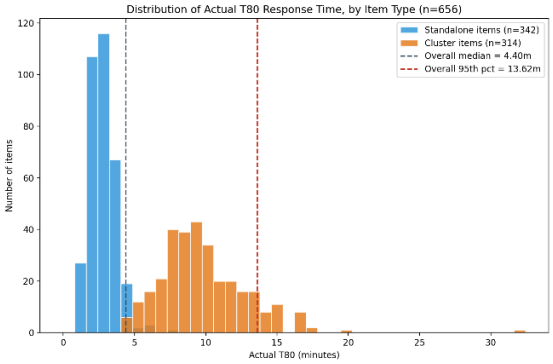


Figure 1. Response Time Distribution

A pre-trained language model "*electra-small-discriminator"* (Clark, Luong, Le & Manning, 2020) was used as the frozen feature extractor to produce informative token representations. With consideration of the limited sample size (N = 656), the low-rank adaptation (LoRA; Hu et al., 2022) was adopted to further fine-tune LLM efficiently. LoRA allows the encoder to adapt at every layer, while keeping the total number of trained parameters small. Two small trainable matrices, A and B, are introduced such that each layer's output is computed as:

$$h = W * x + \left(\frac{\alpha}{r}\right) * B\left(A(x)\right) \qquad (1)$$

where W is the original, pre-trained weight matrix inherited from the frozen language model; A is the first trainable LoRA matrix that projects the input down to a smaller low-rank subspace; B is the second LoRA matrix that projects the low-rank representation produced by A back into the original hidden dimension; r is the LoRA rank and α is a scaling factor.

The proposed model used a two-tier hierarchical attention mechanism (Figure 2) to estimate both the word-level and sentence-level feature importance. At the word level (Tier 1), each sentence was first encoded through *ELECTRA-small,* whose frozen weights were augmented with trainable LoRA matrices injected into the query and value projections of its 12 transformer layers. This yielded a contextualized embedding for each token with a dimension of 256. A learned projection maps to a 64-dimension attention space, from which a single scalar attention score is computed per token. These scores were then normalized via softmax within each sentence, consequently creating a set of word attention weights that sum to one per sentence. The sentence-level representation was then created as the weighted average of the original token embeddings (256 dimensions), using the word attention weights, which yield a distributed sentence representation across sentences.

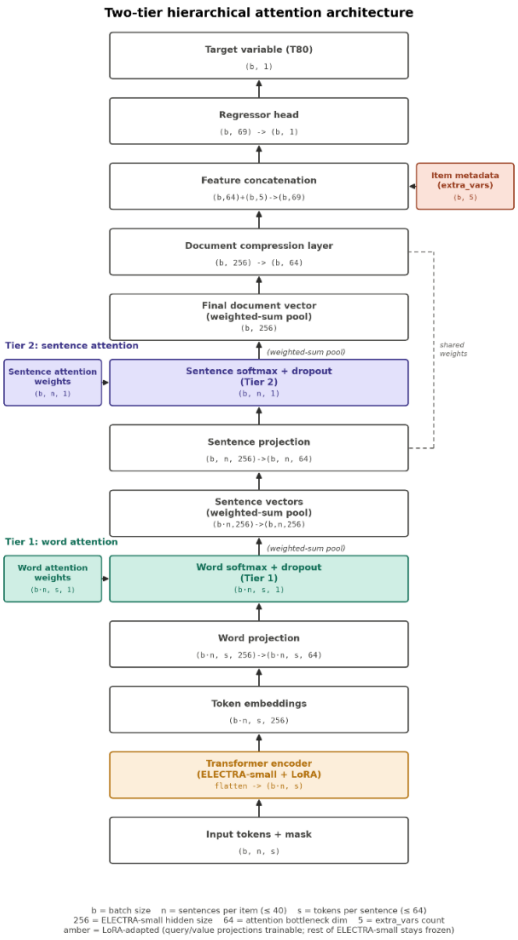


Figure 2. Two-Tier Attention Mechanism

At the sentence level (Tier 2), the sentence representations were treated analogously to the token embeddings of Tier 1, where each sentence representation was projected into a lower-dimensional attention space of 64, and passed through a scoring layer to produce one scalar attention score per sentence. Sentence representations were then pooled into document level as weighted average using sentence attention weights. Before the final regression head, the item attributes were concatenated with the compressed text representation to complement the textual information for the final prediction.

The two tiers were integrated sequentially rather than through an additional learned combination layer. This hierarchical structure allowed the model to learn which words within a sentence, and which sentences within the item text, were most predictive of the target response-time outcome, while retaining information from all words and sentences at every stage of the computation. This serves as a key mechanism for highlighting important text for human review.

2.1 Baseline Models

The Flat Attention Baseline Model: The choice of a two-tier attention architecture as the primary model over a single flat attention mechanism is motivated by three key considerations. First, item text possesses an inherent multi-level structure; explicitly modeling both within-sentence word importance and across-sentence importance allows the network to represent this structure directly, rather than requiring it to be discovered implicitly by an undifferentiated attention layer. Second, the hierarchical design yields more interpretable output: the model produces both a sentence-level attention distribution and, within each sentence, a word-level distribution. This enables qualitative analyses that connect specific predictions to distinct sentence-and-word combinations within the item text. Third, because pre-trained language models have a fixed maximum input length, longer texts must be partitioned before encoding. The hierarchical model performs this partitioning at natural sentence boundaries, whereas a flat model cuts off text at arbitrary token limits, potentially splitting coherent semantic units across independently encoded chunks.

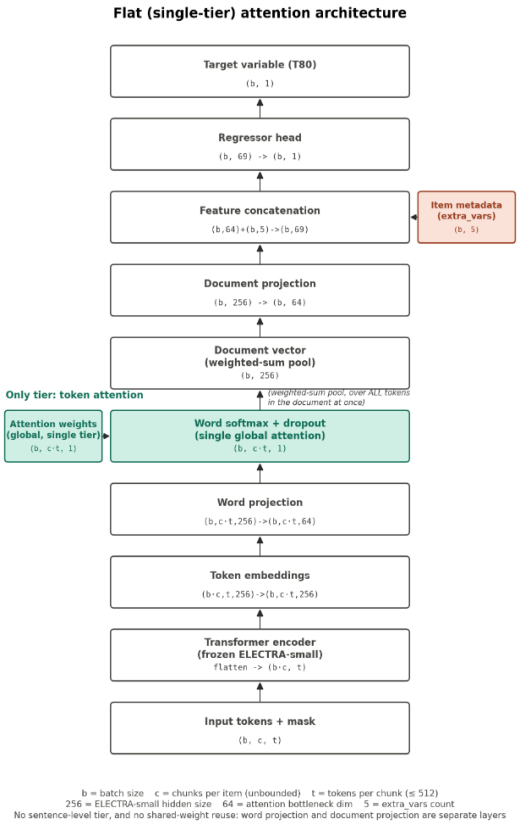


Figure 3. Flat Attention Baseline Model

To empirically evaluate these theoretical advantages, a flat single-level attention architecture was implemented as a primary baseline (Figure 3). This benchmark computes item-level predictions using a single attention layer applied over all tokens in the item text using the frozen *ELECTRA-small* embeddings. Because the encoder enforces a maximum sequence length of 512 tokens, the item text is divided into consecutive, fixed-length token chunks prior to modeling. Each chunk is encoded independently and then concatenating back into a single sequence spanning the full item text. A learned projection maps each token embedding in the reassembled sequence into a 64-dimensional space, where scalar attention scores are calculated and normalized through softmax across the entire document. The weighted sum of the token embeddings produces an aggregate text representation, which is projected into the same lower-dimensional space, concatenated with the item attribute-features, and passed through the regression head to produce the final prediction.

XGBoost Baseline Model: To establish a benchmark evaluating the predictive capacity of the item attributes alone, a gradient-boosted tree model (XGBoost) was trained using only the five non-text item attributes as predictors: figure count, table count, equation count, score point count, and item type. Hyperparameters were fixed based on common defaults for small tabular datasets to achieve a reasonable baseline (Probst et al., 2019).

2.2 Model Training and Evaluation

Both the numeric item attribute variables and the target T_{80} values were normalized prior to training. Model performance was evaluated using Root Mean Squared Error (RMSE) and Mean Absolute Error (MAE) as primary metrics of predictive accuracy.

Both the two-tier and flat attention models were trained using the AdamW optimizer with a mean squared error (MSE) loss function. All models were implemented with PyTorch on the GPU and initialized with a fixed random seed to ensure exact reproducibility across experiments.
To ensure generalizability of the calibrated, a nested cross-validation method was used to separate hyperparameter search from performance estimation. The full dataset (N = 656) was partitioned into 5 folds, with 1 fold being held out as the outer test set. For the inner loop, a further 80/20 split was used for training and validation purposes. Learning rates and epochs were learned through the training set with early stopping on the inner validation loss. The best learning rate and epoch count associated with it were retained to be tested on the outer fold, yielding that fold's MAE and RMSE. This procedure was repeated for all five folds, and average statistics were reported as an estimate of the model's expected performance on unseen items. A simple 5-fold cross-validation split was used to evaluate the XGBoost model.

Hyperparameter Tuning: To balance between model capacity and the risk of overfitting on such small sample, LoRA rank varied across r = 0 (i.e., a fully frozen ELECTRA-

small encoder, with no adapter matrices), r = 4, and r = 8.

A subset of hyperparameters within the two attention-based models were tuned to adapt to the specific characteristics of the item sample (Table 1), with tuning effort concentrated on learning rate and training duration (epochs) while others were held constant. All models were trained for a maximum of 70 epochs, with an early stopping rule applied if validation loss failed to improve for 10 consecutive epochs.

Model	Learning Rate	Epoch	Parameters Fixed (Not Searched)
Flat (1-level attention)	2e-3, 1e-3, 5e-4, 2e-4, 1e-4	70 (via early stopping, patience=10)	maximum token chunk size = 512 per device train batch size = 16 per device eval batch size = 4 weight decay 0.01
Two-tier frozen ELECTRA (r = 0)	2e-3, 1e-3, 5e-4, 2e-4, 1e-4	70 (via early stopping, patience=10)	maximum sentence = 40 maximum words = 64 per device train batch size = 16 per device eval batch size = 4 weight decay 0.01 lora r 0
Two-tier ELECTRA + LoRA (r=4)	2e-3, 1e-3, 5e-4, 2e-4, 1e-4	70 (via early stopping, patience=10)	maximum sentence = 40 maximum words = 64 per device train batch size = 16 per device eval batch size = 4 weight decay 0.01 lora_r 4 lora_alpha 8 lora_dropout 0.05
Two-tier ELECTRA + LoRA (r=8)	2e-3, 1e-3, 5e-4, 2e-4, 1e-4	70 (via early stopping, patience=10)	maximum sentence = 40 maximum words = 64 per_device_train_batch_size = 16 per_device_eval_batch_size = 4 weight_decay 0.01 lora_r 8 lora_alpha 16 lora_dropout 0.05

Table 1. Hyperparameters Fine-Tune by Model

3. Results

The fine-tuned learning rates and epochs per fold are presented in Table 2. Overall, the highest adopted learning rate was 0.0005 between the flat and two-tier frozen attention models, while 0.0010 became the more frequently selected value in LoRA rank 4 and 8 models. Overall, XGBoost obtained the lowest errors across all models, with MAEs below 0.2 in most folds and RMSE below 0.3 in every fold. The flat attention model ranked second-best overall, with lower MAE and RMSE values than any of the two-tier attention model variants. The flat attention model also required the most epochs to train, while the

two-tier attention variants stayed below 21 epochs on average.

Overall, the accuracy scales inversely with model complexity. The more complexity and parameters added to the model, the worse the performance became. It strongly suggests that the larger models tend to overfit with small sample size. It is also shown in the fluctuation in the accuracy indices across folds. For example, in the two-tier attention model with LoRA rank of 8, the MAEs range from 0.2893 to 04954 and RMSE going rom 0.3589 to 0.5630, which indicates that extra model capacity is translating into instability rather than better fit.

Model	Outer Fold	Selected LR	Selected Epochs (Early Stop)	Fold MAE (log-space)	Fold RMSE (log-space)
Flat Attention Model	1	0.0005	44	0.2693	0.3355
	2	0.0005	35	0.2383	0.3086
	3	0.0005	29	0.2328	0.3005
	4	0.0002	69	0.1910	0.2534
	5	0.0005	38	0.2408	0.3117
Average			**43.0**	**0.2345**	**0.3019**
Two-tier Attention Model (frozen)	1	0.0005	34	0.2685	0.3276
	2	0.0005	39	0.3472	0.4270
	3	0.0005	37	0.2394	0.3085
	4	0.0005	31	0.2722	0.3361
	5	0.0005	30	0.2604	0.3248
Average			**34.2**	**0.2775**	**0.3448**
Two-tier Attention Model (LoRA Rank = 4)	1	0.0010	15	0.2677	0.3416
	2	0.0010	14	0.2758	0.3449
	3	0.0005	27	0.2797	0.3427
	4	0.0005	18	0.3594	0.4230
	5	0.0010	18	0.3259	0.4078
Average			**18.4**	**0.3017**	**0.3720**
Two-tier Attention Model (LoRA Rank = 8)	1	0.0005	38	0.2893	0.3715
	2	0.0010	15	0.4954	0.5630
	3	0.0010	10	0.3242	0.3936
	4	0.0005	20	0.3002	0.3589
	5	0.0010	21	0.3201	0.4044
Average			**20.8**	**0.3458**	**0.4183**
XGBoost (attribute-only)	1	N/A	N/A	0.2159	0.2793
	2	N/A	N/A	0.1774	0.2312
	3	N/A	N/A	0.1759	0.2183
	4	N/A	N/A	0.1931	0.2497
	5	N/A	N/A	0.1741	0.2272

Average				**0.1873**	**0.2411**

Table 2. Model Performance by Fold

Figure 4 shows the scatterplots of predicted and actual 80th percentile time across 5 models. The XGBoost's prediction show a distinct gap between 4 to 7 minutes, which mirrors the bimodality of the actual response time shown in Figure 1. The attention-based models show a more continuous, blended spread of predictions across the same region, likely because these models only have indirect access to the item type through text rather than as an explicit feature. Among the attention-based models, the flat attention-based model was the most symmetric in terms of over- and under-estimates. Prediction error increases for all models a actual response time grows, and this effect is more pronounced for the attention-based models than for XGBoost. Notably, the LoRA rank-8 model shows a more severe tendency to overestimate response time than other models.

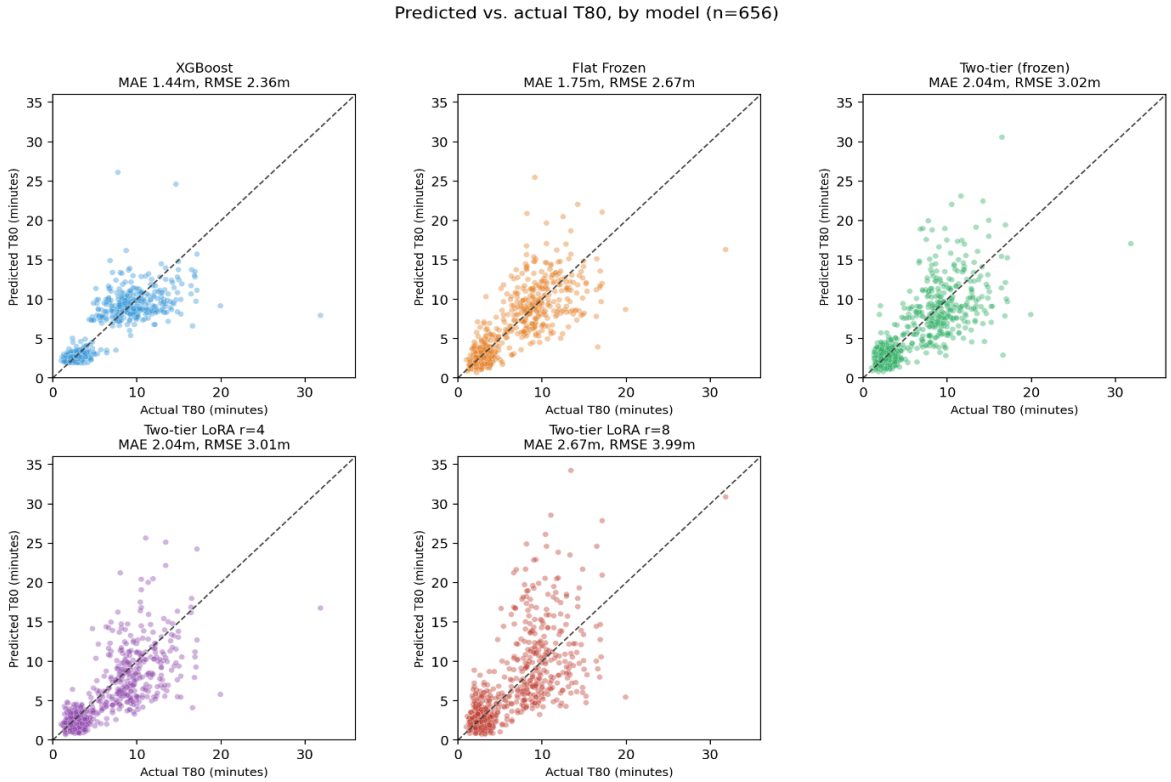


Figure 4. Scatterplots of Predicted and Actual 80th Percentile Response Time

Table 3 reports the weight distribution statistics and word type percentages. The weights are pooled across all words produced by each model. Because the flat model spreads attention across entire item while the two-tier models spread attention within one sentence at a time, the resulting word-level weights cannot be directly compared. We examined each item's single highest-weighted word, identified by the flat and two-iter attention-based models, to assess whether the item's structural characteristics were being captured through attention. If the identified word contains no letters or digits at all, it is counted as punctuation. If it is a noun, verb, or

adjective, it is counted as a content word. Anything else is considered a function word. The flat attention-based model had the highest percentage of all on punctuation (33.2%), while the two-tier model with LoRA rank 4 had the lowest (7.3%). For all models, content words were the most common and punctuation words were the least common, indicating that the models functioned as expected to put emphasis on meaningful words. Once LoRA was applied, the attention shifted markedly toward content words and function words, and punctuation dropped. LoRA r = 4 achieved the lowest punctuation share of all models, while LoRA r = 8 showed the largest increase in content-word attention.

Model	Mean Word-level Weight	SD	% Punctuation	% Function Word	% Content Word
Flat (frozen)	0.00265	0.00292	33.2	16.0	50.8
Two-tier (frozen)	0.04761	0.03804	14.5	23.8	61.7
Two-tier LoRA r=4	0.05000	0.04470	7.3	27.4	65.2
Two-tier LoRA r=8	0.04762	0.04301	9.0	12.7	78.4

Table 3. Word Level Weights Statistics and Word Type Importance Percentages

Taken together, these findings support the paper's central framing: the two-tier attention architecture's value lies in interpretability. Adapting the encoder with LoRA further transformed what attention identifies as important, without sacrificing predictive accuracy at r = 4. This suggests that even minimal, parameter-efficient adaptation of the encoder is sufficient to redirect the model's attention towards linguistically meaningful content.

The primary strength of the proposed model lies in its capacity for exporting word- and sentence-level weights that can map back onto the raw item text for human interpretability. These weights serve as diagnostic output, which can be rendered into a highly accessible format. By visually highlighting the core friction words in a sentence, the model provides granular, text-based insights. This interpretative capability creates a continuous feedback loop that enables content writers to engage in—making targeted adjustments and immediate evaluation of the predicted impact on response time.

4. Conclusions and Future Work

In this paper, the findings of this study demonstrate the viability of using the two-tier attention architecture to model student response times and provide interpretable diagnostic feedback for item development. To build upon these results, future research should expand the item sample size across a broader range of academic domains, such as English Language Arts assessments, where reading comprehension drives the majority of response latency.

Two limitations motivate future work with a larger item sample. First, LoRA rank 8 did not improve on rank 4 — it showed greater fold-to-fold instability and lower accuracy, indicating that this dataset supports only a small amount of additional encoder capacity before overfitting risk outweighs any benefit. Second, sentence-level attention remained close to uniform across sentences within an item, indicating limited ability to discriminate important sentences from unimportant ones; this will also need further investigation with a larger sample. Additionally, future iterations of this work will focus on integrating true multimodal architectures capable of directly processing embedded figures, tables, and mathematical equations, rather than relying on scalar counts of these elements as proxy features.

5. Acknowledgements

I thank my colleagues Claire Chang, Suhwa Han, Zebing Wu, and Mina Lee for their earlier work on the topic and helpful discussions around this work.

References

Ackerman, P. L., & Kanfer, R. (2009). Test length and cognitive fatigue: an empirical examination of effects on performance and test-taker reactions. *Journal of Experimental Psychology: Applied*, *15*(2), 163. https://doi.org/10.1037/a0015719Phillip L. Ackerman and Ruth Kanfer. 2009. Test length and cognitive fatigue: An empirical examination of effects on performance and test-taker reactions. *Journal of Experimental Psychology: Applied*, 15(2):163–181. https://doi.org/10.1037/a0015719

Baldwin, P., Yaneva, V., Mee, J., Clauser, B. E., & Ha, L. A. (2021). Using natural language processing to predict item response times and improve test construction. *Journal of Educational Measurement*, *58*(1), 4-30. Peter Baldwin, Victoria Yaneva, John Mee, Brian E. Clauser, and Le An Ha. 2021. Using natural language processing to predict item response times and improve test construction. *Journal of Educational Measurement*, 58(1):4–30. https://doi.org/10.1111/jedm.12264

Brent Bridgeman and Frederick Cline. 2000. Variations in mean response times for questions on the computer-adaptive GRE® general test: Implications for fair assessment. *ETS Research Report Series*, 2000(1):i–29. https://doi.org/10.1002/j.2333-8504.2000.tb01830.x

Bridgeman, B., & Cline, F. (2000). Variations in mean response times for questions on the computer-adaptive GRE® general test: implications for fair assessment. *ETS Research Report Series*, *2000*(1), i-29.

Clark, K., Luong, M., Le, Q. V., & Manning, C. D. (2020). ELECTRA: Pre-training text encoders as discriminators rather than generators. *Proceedings of the International Conference on Learning Representations (ICLR)*. https://doi.org/10.48550/arXiv.2003.10555

Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. LoRA: Low-Rank Adaptation of Large Language Models. In *Proceedings of the International Conference on Learning Representations (ICLR)*.

Hariram Veeramani, Surendrabikram Thapa, Natarajan Balaji Shankar, and Abeer Alwan. 2024. Large Language Model-based Pipeline for Item Difficulty and Response Time Estimation for Educational Assessments. In *Proceedings of the 19th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2024)*, pages 561–566, Mexico City, Mexico. Association for Computational Linguistics. https://aclanthology.org/2024.bea-1.49/

Kevin Clark, Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. ELECTRA: Pre-training text encoders as discriminators rather than generators. In *Proceedings of the 8th International Conference on Learning Representations (ICLR)*. https://doi.org/10.48550/arXiv.2003.10555

Peter E. D. Love, Weili Fang, Jane Matthews, Stuart Porter, Hanbin Luo, and Lieyun Ding. 2023. Explainable artificial intelligence (XAI): Precepts, models, and opportunities for research in construction. *Advanced Engineering Informatics*, 57:102024.

Philipp Probst, Bernd Bischl, and Anne-Laure Boulesteix. 2019. Tunability: Importance of hyperparameters of machine learning algorithms. *Journal of Machine Learning Research*, 20(53):1–32.

Veeramani, H., Thapa, S., Shankar, N. B., & Alwan, A. (2024). Large language model-based pipeline for item difficulty and response time estimation for educational assessments. *Proceedings of the 19th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2024)* (pp. 561-566). https://aclanthology.org/2024.bea-1.49/

Yang, Z., Yang, D., Dyer, C., He, X., Smola, A., & Hovy, E. (2016). Hierarchical attention networks for document classification. *Proceedings of the 2016 conference of the North American chapter of the association for computational linguistics: human language technologies* (pp. 1480-1489). San Diego, California. Association for Computational Linguistics. https://doi.org/10.18653/v1/N16-1174

Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alex Smola, and Eduard Hovy. 2016. Hierarchical Attention Networks for Document Classification. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for*

Computational Linguistics: Human Language Technologies, pages 1480–1489, San Diego, California. Association for Computational Linguistics. https://doi.org/10.18653/v1/N16-1174

Beyond Agreement: Calibrating Automated Scoring to Human Judgment Using Control Scripts

Mark Dulhunty, Alex Codoreanu, and Nathan Zoanetti
Australian Council for Educational Research
Melbourne, Victoria, Australia

Correspondence: mark.dulhunty@acer.org

Abstract

Automated scoring systems are typically evaluated using agreement with human-assigned scores, yet deployment in operational assessment requires evidence of calibration, reproducibility, and alignment with existing quality assurance processes. This study evaluates an ensemble automated scoring approach across 55 constructed-response reading items using approximately 260,000 student responses. Initial large-scale evaluation established strong agreement with human-assigned scores. Model performance was subsequently calibrated against operational reference scores using 477 embedded control scripts and compared with 27 human markers. The ensemble achieved 95.3% agreement with operational reference scores, matching or exceeding individual human marker performance on many items. Reproducibility testing demonstrated complete scoring stability across repeated evaluations, providing evidence that the approach generated consistent outcomes under operational conditions. High-performing models were then applied within marking workflows to support quality assurance through targeted review of disagreements between automated and human-assigned scores. The findings demonstrate how automated scoring can be evaluated and calibrated within operational assessment systems, supporting its use as an embedded measurement and quality assurance function rather than solely a predictive tool.

1 Introduction

Evaluating automated scoring in operational assessment requires situating model performance relative to both scoring standards and observed human marker behaviour (Frohn et al., 2025b; Melconian, 2026; Rotou & Rupp, 2020). Much existing research, including recent large language model and ensemble-based approaches (e.g., Frohn et al., 2025a; Xue et al., 2026), treats automated scoring as a predictive task, emphasising accuracy relative to reference scores (e.g., Ahmed et al., 2022; Henkel et al., 2025; Jung et al., 2026). In operational contexts, however, it must function as a measurement instrument, requiring evaluation beyond accuracy to consider alignment with human judgement and quality assurance practices.

This study addresses this gap through a multi-phase, program-aligned evaluation, of automated scoring in practice. An initial large-scale evaluation was conducted to establish model performance relative to human markers (N ≈ 260,000 responses), providing evidence of reliability at scale. Building on this, control scripts (i.e., responses with established reference scores) were embedded within operational marking. Throughout this paper, reference scores are treated as the operational scoring standard. These scores were established through lead-marker consensus and should be interpreted as operationally validated rather than objectively 'true' scores (see Sample section for details). These scripts were scored by both human markers and automated models, enabling evaluation against reference standards and direct comparison with individual human marker performance. In summary, the study follows a three-phase evaluation logic: large-scale evaluation to establish reliability, control script evaluation against reference scores, and operational deployment to assess practical utility.

2 Aims

This study aims to:

1. Evaluate agreement of the ensemble automated scoring approach relative to

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 79–85
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

human scoring at scale across constructed-response items.
2. Assess calibration through comparison with reference scores and individual human marker performance on control scripts.
3. Examine the extent to which the model can be applied to operational responses to support identification of inconsistent human judgements.

3 Sample(s)

The study draws on multi-phase operational assessment data from a large-scale reading program. The primary evaluation dataset comprises 55 dichotomously scored constructed-response items and approximately 260,000 student responses. Responses were typically short (ranging from phrases to one or two sentences) such as "*There is a supermarket that takes up a lot of room*".

Control scripts (N = 477) with established reference scores were embedded within operational marking and scored by at least two lead markers, with individual markers (N = 27) scoring varying subsets of scripts due to operational allocation constraints. The number of control scripts varied by item (ranging from 4 to 14 control scripts depending on the item).

A subset of items (N=10), for which the ensemble automated scoring approach demonstrated strong performance, was selected for a later operational phase (N ≈ 14,000 responses) to evaluate its application for quality assurance.

4 Methods

The study follows three phases: large-scale evaluation to establish reliability, validation against control scripts with known scores, and operational deployment to assess practical utility.

Control scripts with established reference scores were embedded within operational marking, allowing human and automated scoring of shared responses under operational conditions. This design enabled direct comparison of model outputs with individual human markers within a shared measurement framework.

Automated scoring combined generative and embedding-based approaches within an ensemble framework (see Appendix A for full model specifications and Table A1). The ensemble integrated outputs from five model components differing in architecture, training data, and prompt configuration, with final scores assigned using a majority voting rule. Although components were not fully independent due to shared data sources, this design introduces controlled diversity, improving robustness and consistency (Xue et al., 2026).

Agreement estimates were calculated using approximately 240,000 held-out responses, with approximately 10% of available responses used for model training across the transformer based ensemble components. Model performance was evaluated against reference scores and relative to individual human marker performance on control scripts, supporting comparison against both scoring standards and observed variability in human marker behaviour. Additional agreement analyses were conducted on operational responses using human-assigned scores to assess broader performance beyond control scripts.

To assess reproducibility, a subset of approximately 50,000 responses was repeatedly scored using the same ensemble configuration. This analysis was conducted to evaluate score stability and address potential concerns that generative LLM components may produce different outputs when repeatedly presented with the same response. Repeated scoring outcomes were compared to assess the consistency of score assignments across runs.

Following this evaluation, a subset of items was selected for which the ensemble automated scoring approach demonstrated strong performance based on agreement and calibration relative to both reference scores and human marker performance on control scripts. Selection was further informed by consistently high agreement, defined as exceeding 95% agreement during initial evaluation across all individual scoring components. The model was then applied to operational responses for these items to generate disagreement reports during marking, comparing automated and human-assigned scores to identify responses for targeted review within existing quality assurance processes.

The overall evaluation and deployment framework is illustrated in Figure 1.

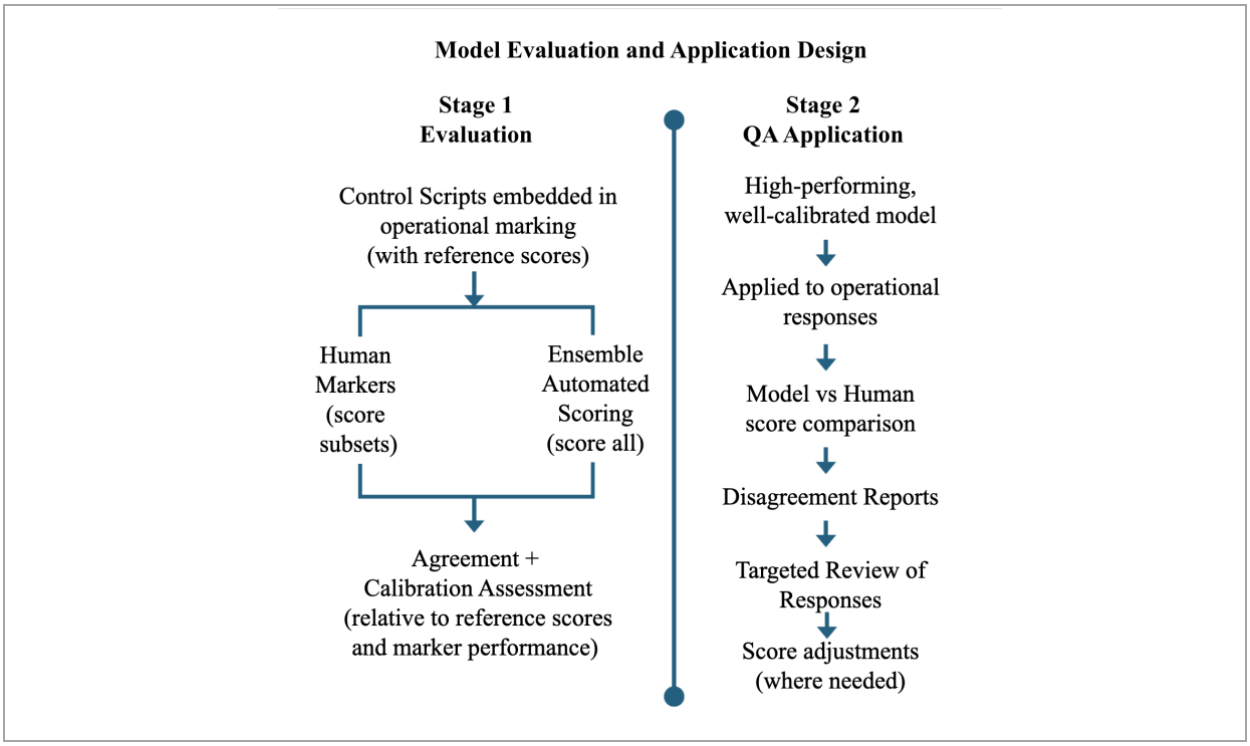


Figure 1: Following large-scale evaluation, embedded control scripts enable human-referenced evaluation of automated scoring (Stage 1), supporting deployment of calibrated models for quality assurance (QA) through disagreement detection (Stage 2).

5 Results

Predicted scores for each component of the ensemble automated scoring approach were initially evaluated against approximately 260,000 human-marked responses across 55 items. Each component achieved an average item-level Cohen's kappa (κ) greater than 0.70 (range: 0.73–0.78), indicating strong agreement with human markers and supporting progression to subsequent evaluation phases.

Evaluation of the final ensemble demonstrated similarly strong performance across most items (Figure 2). Of the 55 items evaluated, 44 (80%) achieved item-level κ values above 0.6, including 26 items (47%) above κ = 0.8. Only two items returned κ values below 0.4, with the majority of items demonstrating moderate to strong agreement.

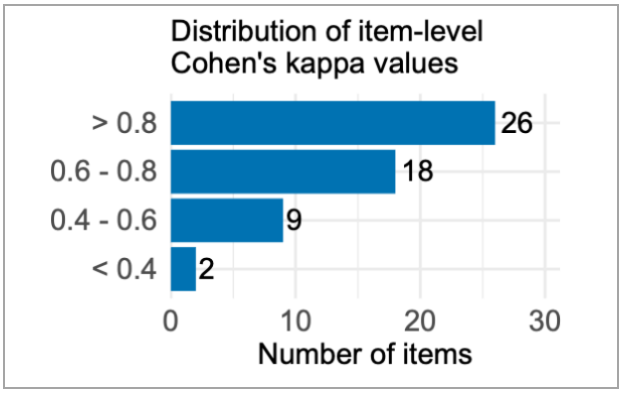


Figure 2: Distribution of item-level Cohen's κ values for the ensemble automated scoring approach across 55 reading constructed-response items.

The ensemble approach was then evaluated against 477 control scripts with established reference scores. Across these control scripts, the ensemble achieved 95.3% agreement with operational reference scores. In addition, perfect agreement was observed for 39 of 55 items (71%), although individual items contained between 4 and 14 control scripts. Performance on many items was comparable to or exceeded that of many individual human markers.

The ensemble 'AutoScore' performance against individual human marker performance on control scripts is shown in Figure 3. Figure 3 shows that AutoScore performed similarly to or better than most markers within the comparison pool.

In Figure 3, 'Similar' performance was when AutoScore was within +/-0.5% of human marker accuracy. Comparisons were restricted to markers with substantial evidence available (≥ 21 items and ≥ 180 control scripts). AutoScore performed similarly to or better than 18 of the 22 markers in the comparison pool. AutoScore underperformed compared to 4 markers, who were also the top performing markers across the entire marking pool.

Direct comparison with individual marker performance showed that model scores were well calibrated relative to both reference standards and observed variability in human marking behaviour.

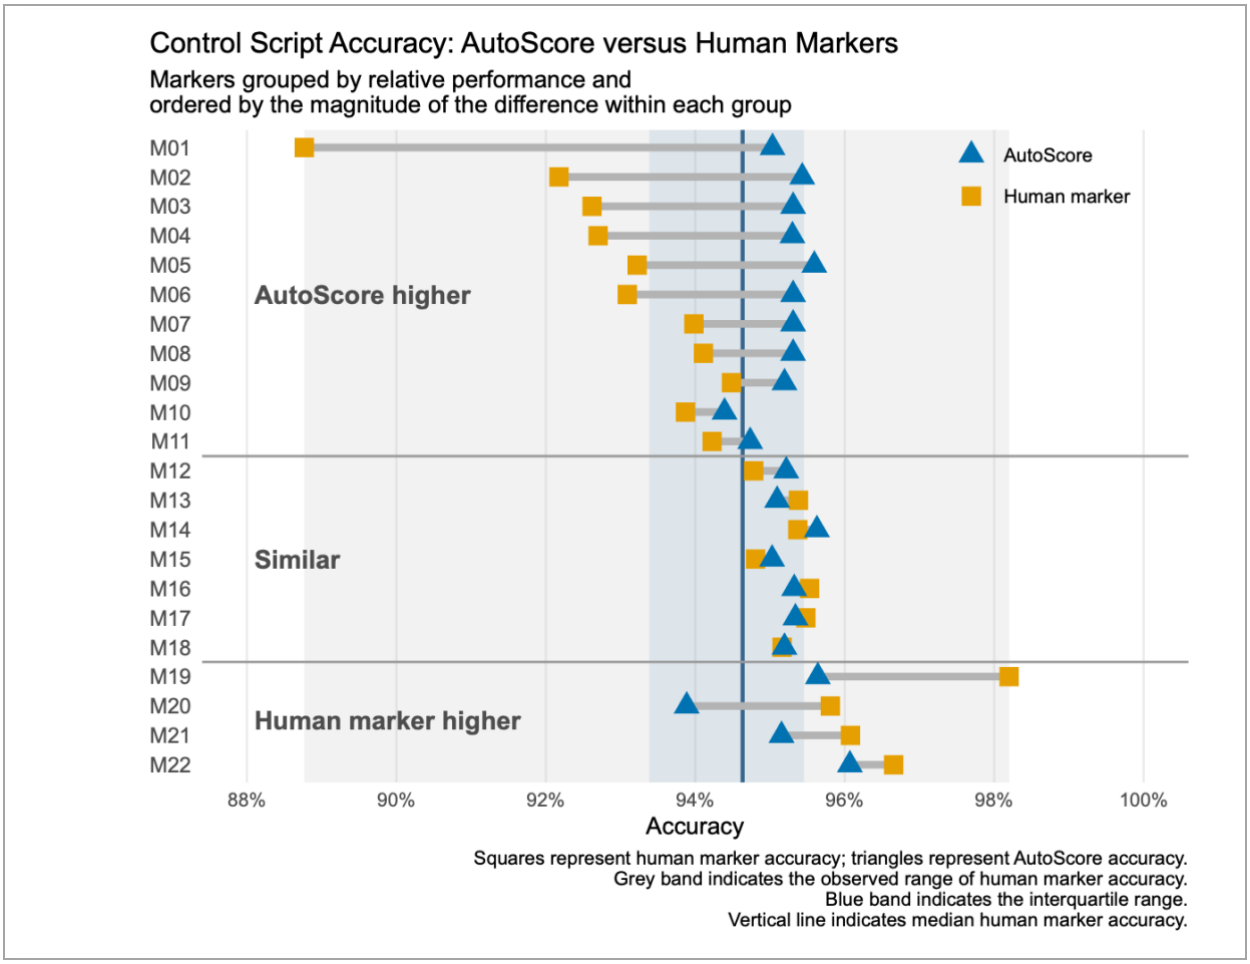


Figure 3: Comparison of AutoScore and human marker accuracy on control scripts. Markers are grouped according to whether AutoScore achieved higher, similar, or lower accuracy than the human marker. Detailed marker-level results are provided in Appendix B.

Reproducibility testing demonstrated complete scoring stability across repeated evaluations, with identical scores assigned when the same responses were evaluated repeatedly (based on a sample of approximately 50,000 responses).

When applied to operational responses, high-performing models identified a small subset of disagreements indicative of potential scoring inconsistencies. Although overall agreement was high and the number of disagreements was correspondingly low, targeted review of these cases resulted in 13 confirmed score adjustments from approximately 14,000 operational responses, demonstrating the capacity of automated scoring to support existing quality assurance processes.

The proportion of disagreements was lower than typical quality assurance sampling thresholds, suggesting that the approach can support more

targeted and efficient review while maintaining confidence in scoring outcomes. Taken together, these findings indicate that well-calibrated automated scoring models can be deployed to support quality assurance processes through the targeted identification of potentially inconsistent human judgements, while preserving existing operational workflows.

6 Conclusions

Findings show that well-calibrated automated scoring models can achieve consistent performance across items, in some cases matching or exceeding individual human marker performance on control scripts while remaining aligned with established scoring standards. When integrated into operational workflows, such models can support quality assurance through targeted review of potentially inconsistent human judgements without disrupting existing processes.

Reliable operational scoring requires that identical responses receive identical scores when processed repeatedly. Reproducibility testing therefore complements agreement and calibration analyses by confirming that the automated scoring process behaves deterministically in practice. Results demonstrated complete scoring stability across repeated evaluations, showing that ensemble-based automated scoring can provide stable and repeatable scoring outcomes despite incorporating generative LLM components.

This study demonstrates that automated scoring can be evaluated and calibrated within operational assessment systems through comparison with both reference scores and observed human marker behaviour on control scripts. This human-referenced calibration framework provides a practical basis for interpreting model performance relative to operational scoring standards, extending evaluation beyond aggregate accuracy metrics alone.

More broadly, this work reframes automated scoring from a purely predictive task to an embedded measurement function, demonstrating how model evaluation, calibration, and deployment can be integrated within ongoing scoring and quality assurance practices.

This study forms part of a broader programme of research investigating automated scoring methods, infrastructure, and operational deployment within large-scale assessment systems. Future work will extend the calibration framework presented here to additional item types, broader rubric-based scoring approaches, and other applications of AI-assisted assessment.

7 Limitations

This study has several limitations. First, the number of control scripts per item was relatively small (4-14), reflecting operational rather than experimental data collection. Second, control scripts were selected for quality assurance purposes rather than randomly sampled, which may limit generalisability. Third, individual markers completed differing subsets of control scripts because of operational allocation requirements, resulting in unequal evidence across marker-level comparisons. Finally, the study focused on dichotomously scored reading items within a single assessment program, and findings may not generalise to other item types or scoring contexts.

Acknowledgements

The authors gratefully acknowledge the support of ACER colleagues, particularly Lynn Sendy-Smithers, Jenna Schorel, and Rose Knight, whose contributions helped make this work possible.

References

Ahmed, A., Joorabchi, A., & Hayes, M. J. (2022). On the application of sentence transformers to automatic short answer grading in blended assessment. In *2022 33rd Irish Signals and Systems Conference (ISSC 2022)*. Institute of Electrical and Electronics Engineers. https://doi.org/10.1109/ISSC55427.2022.9826194

Frohn, S., Burleigh, T., & Chen, J. (2025a). Automated scoring of short answer questions with large language models: Impacts of model, item, and rubric design. In A. I. Cristea, E. Walker, Y. Lu, O. C. Santos, & S. Isotani (Eds.), *Artificial Intelligence in Education (AIED 2025)* (Lecture Notes in Computer Science, Vol. 15882). Springer. https://doi.org/10.1007/978-3-031-98465-5_6

Frohn, S., Deters, L., Burleigh, T., & Dicerbo, K. (2025b). A framework for live scoring constructed response items with commercial LLMs. Paper presented at the Artificial Intelligence in Measurement and Education Conference, Pittsburgh, PA. https://www.researchgate.net/publication/396937943_A_Framework_for_Live_Scoring_Constructed_Response_Items_with_Commercial_LLMs

Henkel, O., Hills, L., Roberts, B., et al. (2025). Can LLMs grade open response reading comprehension questions? An empirical study using the ROARs dataset. *International Journal of Artificial Intelligence in Education, 35*, 651–676. https://doi.org/10.1007/s40593-024-00431-z

Jung, J. Y., Bezirhan, U., & von Davier, M. (2026). Optimizing automated scoring in ILSAs with prompt compression. *Computers & Education: Artificial Intelligence*. https://doi.org/10.1016/j.caeai.2026.100558

Melconian, M. (2026, April 8). Inter-rater agreement and AI scoring: The reliability evidence you now need. LinkedIn. https://www.linkedin.com/pulse/inter-rater-agreement-ai-scoring-reliability-evidence-melconian-7cpme

Qwen Team. (2024, September). *Qwen2.5: A party of foundation models*. https://qwenlm.github.io/blog/qwen2.5/

Reimers, N., & Gurevych, I. (2019). Sentence-BERT: Sentence embeddings using Siamese BERT-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. https://arxiv.org/abs/1908.10084

Rotou, O., & Rupp, A. A. (2020). Evaluations of automated scoring systems in practice (ETS Research Report No. RR-20-10). *ETS Research Report Series*. https://doi.org/10.1002/ets2.12293

Sentence Transformers. (n.d.-a). *all-mpnet-base-v2*. Hugging Face. https://huggingface.co/sentence-transformers/all-mpnet-base-v2

Sentence Transformers. (n.d.-b). *all-roberta-large-v1*. Hugging Face. https://huggingface.co/sentence-transformers/all-roberta-large-v1

Tunstall, L., Reimers, N., Jo, U. E., Bates, L., Korat, D., Wasserblat, M., & Pereg, O. (2022). Efficient few-shot learning without prompts: SetFit. *arXiv*. https://arxiv.org/abs/2209.11055

Xue, M., Xiao, X., Liu, Y., & Wilson, M. (2026). On the consistency of automatic scoring with large language models. *Educational and Psychological Measurement*. https://doi.org/10.1177/00131644261418138

Appendix A Model Specifications

Automated scoring combined generative and embedding-based approaches. A large language model (Qwen2.5-14B; Qwen Team, 2024) was used for direct scoring, while sentence-transformer models (all-mpnet-base-v2 and all-roberta-large-v1; Sentence Transformers, n.d.-a, n.d.-b) were fine-tuned using SetFit (Tunstall et al., 2022) to learn task-specific scoring functions from labelled responses, building on the Sentence-BERT framework (Reimers & Gurevych, 2019). Model configurations, training sample sizes, and prompt conditions are summarised in Table A1.

Setting	Model	Approach Type	Training Samples	Training Sample Set	Prompt / Input Configuration
1	all-mpnet-base-v2	Sentence Transformer	100	Sample 1	Encoded response only
2	all-mpnet-base-v2	Sentence Transformer	125	Sample 2	Encoded response only
3	all-roberta-large-v1	Sentence Transformer	125	Sample 2	Encoded response only
4	Qwen2.5-14B	LLM	5	Sample 3	Question stem + scoring criteria
5	Qwen2.5-14B	LLM	5	Sample 3	Question stem + scoring criteria + **stimulus**

Table A1: Model configurations included in the ensemble automated scoring approach, showing variation in model type, training data, and prompt design across components.[1]

[1] Note. Sentence transformer models were fine-tuned using SetFit. "Training Samples" refers to number of samples per score category. "Training Sample Set" denotes distinct subsets of audited responses, with some overlap across configurations as described in the Methods section.

Appendix B Detailed Marker Results

Marker Label	Number of Items	Total Scripts	Mean Human Accuracy	Mean AutoScore Accuracy	Difference
M01	30	296	88.8%	95.0%	6.3%
M02	51	457	92.2%	95.4%	3.3%
M03	55	477	92.6%	95.3%	2.7%
M04	40	387	92.7%	95.3%	2.6%
M05	54	472	93.2%	95.6%	2.4%
M06	55	477	93.1%	95.3%	2.2%
M07	55	477	94.0%	95.3%	1.3%
M08	55	477	94.1%	95.3%	1.2%
M09	31	301	94.5%	95.2%	0.7%
M10	23	229	93.9%	94.4%	0.5%
M11	49	432	94.2%	94.7%	0.5%
M12	23	242	94.8%	95.2%	0.4%
M13	21	195	95.4%	95.1%	-0.3%
M14	43	418	95.4%	95.6%	0.3%
M15	36	348	94.8%	95.0%	0.2%
M16	34	338	95.5%	95.3%	-0.2%
M17	50	445	95.5%	95.3%	-0.1%
M18	31	301	95.2%	95.2%	0.0%
M19	25	233	98.2%	95.6%	-2.6%
M20	26	181	95.8%	93.9%	-1.9%
M21	26	180	96.1%	95.2%	-0.9%
M22	29	266	96.7%	96.1%	-0.6%

Table B1: Human marker control script scoring accuracy details with accompanying AutoScore comparison.

Human-in-the-Loop Gemini Item Generation for Large-Scale MCQ Banks

Melchor Sánchez-Mendiola, Milton García-Lima, Elibidú Ortega-Sánchez, Manuel García-Minjares, and Enrique Buzo-Casanova

Coordination of Evaluation, Innovation and Educational Development (CEIDE), National Autonomous University of Mexico (UNAM), Mexico City, Mexico

melchorsm@unam.mx; milton_garcia@ceide.unam.mx; elibidu_ortega@ceide.unam.mx; manuel_garcia@ceide.unam.mx; enrique_buzo@ceide.unam.mx

Abstract

Large-scale assessment programs need secure, content-aligned item banks, yet faculty item writing is resource intensive and variable. We report a work-in-progress, naturalistic comparison of 2,405 Spanish-language multiple-choice questions developed at the National Autonomous University of Mexico: 986 drafted with ten customized Gemini Gems and 1,419 drafted by faculty. A Gem is a reusable Gemini configuration with persistent instructions and optional knowledge files. The subject-specific Gems embedded assessment blueprints, diagnostic gatekeeping, item-writing constraints, and a standardized output schema. Three-member faculty committees reviewed every candidate item using institutional criteria. We summarized review activity with a validation-friction index (VFI; recorded comments divided by submitted items) and analyzed recurring comment patterns. Overall VFI was 1.84 for Gemini-drafted items and 2.30 for faculty-drafted items; the AI source had lower VFI in three of four knowledge areas. Gemini drafts were more consistent in blueprint adherence and structure, whereas faculty drafts often offered richer contextualization and misconception-based distractors. Because comments mixed confirmations and corrections and no examinee data were available, VFI is a process indicator, not evidence of psychometric equivalence. Findings support a human-in-the-loop workflow combining structured generation, automated checks, expert review, and post-administration validation.

1 Introduction

Large-scale assessment programs depend on item banks that are sufficiently broad, secure, and aligned with intended content and cognitive demands. Producing those banks requires much more than subject-matter expertise, each multiple-choice question (MCQ) must represent the intended construct, use clear language, contain a defensible key, and include plausible distractors without grammatical, lexical, or visual clues. Weaknesses at this stage can introduce construct-irrelevant variance before an item is ever administered (American Educational Research Association et al., 2014; Haladyna et al., 2002; Haladyna and Rodriguez, 2013).

Traditional faculty authoring supports disciplinary authenticity and can draw on knowledge of common student misconceptions. It is also labor intensive, difficult to standardize across many authors, and challenging to scale when programs need multiple test forms or frequent bank renewal. Earlier automatic item generation (AIG) approaches addressed this problem through cognitive models, item models, templates, and content coding. These approaches demonstrate that human judgment and computational generation can be combined to improve production capacity and test security (Circi et al., 2023; Gierl et al., 2022).

Large language models (LLMs) extend AIG by producing fluent items from natural-language instructions and curricular materials. The resulting flexibility is attractive, but fluency is not equivalent to measurement quality. Reviews and empirical comparisons have identified usable AI-generated items alongside ambiguous keys, weak or implausible distractors, inappropriate difficulty, and content-alignment problems (Artsi et al., 2024; Grévisse et al., 2024; Kaldaras et al., 2024). Recent studies using examinee data have also reported mixed results: AI-generated items may be comparable on some indices yet easier, differently discriminating, or more readily detectable than human-authored items (Kaya et al., 2025; Law et al., 2025; Linde et al., 2026). Across this literature, expert review remains a central safeguard.

This work-in-progress paper examines a different but complementary stage of the validity

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 86–92
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

process: pre-operational item review in a large, cross-disciplinary, Spanish-language item-bank project. We compare items drafted with customized Google Gemini Gems and items written by faculty at the National Autonomous University of Mexico (UNAM). The study asks three questions: How did the two item sources differ in recorded expert-review activity across four knowledge areas? What validity-relevant strengths and weaknesses recurred in reviewer comments? What workflow implications follow for responsible, scalable, human-in-the-loop item development?

2 Institutional Context

UNAM's Coordination of Evaluation, Innovation and Educational Development (CEIDE) developed the project to support a large institutional MCQ bank spanning four knowledge areas: Physical-Mathematical Sciences and Engineering (PMSE), Biological, Chemical, and Health Sciences (BCHS), Social Sciences (SS), and Humanities and Arts (HA). The assessment blueprint specified 120 items per area for each of five versions, or approximately 2,400 items. The realized comparison dataset contained 2,405 candidate items across ten subjects: Mathematics, Spanish, Physics, Chemistry, Biology, Literature, Geography, Philosophy, Mexican History, and World History.

Two authoring routes contributed items to the same institutional development process. Faculty authors wrote items using the applicable table of specifications. In parallel, ten subject-specific Gemini Gems generated drafts from the same types of inputs: subject, topic, expected learning outcome, intended cognitive level, item format, and media requirements. The comparison was therefore a retrospective, naturalistic process evaluation rather than a randomized experiment. Item source was not assigned within matched blueprint cells, and the number of AI- and faculty-drafted items differed by area.

3 Gemini Gem Item-Generation Architecture

Google Gemini 3.0 Pro was used. Google describes a custom Gem as a saved configuration that can retain instructions and optional knowledge files for repeated use (Google, n.d.). In this project, a Gem differed from a one-off prompt in three practical respects: the same instructions were reused across sessions, curricular and item-writing documents could be attached as persistent knowledge, and changes to the central instruction set could be propagated to subsequent generations. The Gems were configured rather than fine-tuned; no new foundation model was trained.

The available project records identify the application as Gemini Gems but do not preserve an auditable foundation-model version identifier. Accordingly, this paper evaluates the configured workflow as implemented, not a frozen-model benchmark. Future replications should log the model label, generation dates, account tier, prompt version, knowledge-file versions, and any accessible inference settings.

The architecture combined four design components. First, each Gem received a subject and assessment-expert persona. Second, a diagnostic gate required the user to specify the relevant blueprint information before generation; the explicit instruction was not to generate an item until the diagnostic phase was complete. Third, the Gems embedded institutional item-writing constraints, including Mexican Spanish, positive wording, one defensible key, balanced options, plausible distractors, structured metadata, and LaTeX-compatible notation for scientific and mathematical content. For database standardization, the draft key was placed in option A and could be randomized after import. Fourth, when an item required a figure, the Gem produced the editable text item and a separate image-generation specification rather than rendering the entire item as an image.

3.1 Prompt Scaffold and Calibration

Table 1 presents an abridged English translation of the documented prompt scaffold. It does not reproduce confidential operational content or any item. The design sought to make curricular grounding and output constraints explicit before the model generated a draft. This is consistent with recommendations that AI-assisted test development document the construct, intended use, prompt design, semantic alignment, and human validation steps (Kaldaras et al., 2024; Suárez-Álvarez et al., 2026).

The architecture also specified an iterative calibration loop. When reviewers identified a recurring methodological or stylistic problem, the project team could revise the central instructions by adding a restriction or a counterexample. The first production audit identified priorities for subsequent prompt revisions, including shorter stems, more symmetric option length, less lexical overlap

between stem and key, and distractors grounded more directly in common misconceptions. Because the available records describe the calibration mechanism but do not contain a complete prompt-version history, the present analysis does not estimate improvement across prompt iterations.

Component	Illustrative translated instruction	Quality-control function
Persona and scope	You are a [subject] instructor and an expert in educational assessment. Work only within the specified curriculum and assessment standards.	Stabilizes disciplinary role, register, and intended task.
Diagnostic gate	Do not generate an item until the user supplies the topic, expected learning outcome, cognitive level, item type, and media requirements.	Reduces premature generation and forces blueprint completion.
Item constraints	Use Mexican Spanish and positive wording; provide one defensible key and three plausible distractors based on common student errors; avoid grammatical, lexical, and visual clues.	Operationalizes institutional item-writing rules.
Technical format	Use a fixed metadata schema; place the draft key in option A for import and later randomization; use LaTeX-compatible notation for mathematics and science.	Supports parsing, database transfer, and reproducible metadata.
Multimodal handling	When a visual is required, return the text item plus a separate, precise image specification rather than rendering the complete item as an image.	Preserves editable item structure and separates content from rendering.
Calibration	When a recurrent flaw is found, add a specific restriction or a counterexample to the central instructions and retain a prompt revision record.	Enables centralized, auditable refinement across future generations.

Table 1: Abridged and translated Gemini Gem instruction scaffold. Bracketed terms were replaced with the applicable subject and blueprint information.

4 Data and Methods

The unit of analysis was a candidate MCQ submitted to the institutional review process. Of the 2,405 items, 986 were drafted with Gemini Gems and 1,419 were drafted by faculty. Items were distributed across the four areas as shown in Table 3. The project records contained item-source counts, the total number of recorded committee comments, and area-level qualitative audits of those comments.

4.1 Faculty Review Process

Every candidate item underwent external content review by a committee of three faculty experts in the relevant subject. Reviewers assessed the item against institutional criteria covering blueprint alignment, cognitive level, clarity, technical correctness, a unique defensible key, distractor plausibility, and absence of unintended clues. Committees recorded one or more comments in the item-management workflow. Importantly, the comment field included both positive confirmations (for example, that an item met a criterion) and requests for revision. Consequently, a comment count cannot be interpreted as an error count.

The review criteria are summarized in Table 2. They functioned as pre-operational content and technical quality checks. They did not establish score reliability, difficulty, discrimination, differential item functioning, or validity evidence based on examinee response processes or relationships with external variables.

Review dimension	Primary question for the faculty committee
Blueprint alignment	Does the item assess the specified subject, topic, and expected learning outcome?
Cognitive demand	Does the item elicit the intended level of knowledge, comprehension, or application?
Clarity and language	Is the stem understandable, concise, grammatically sound, and appropriate in Mexican Spanish?
Key validity	Is there one clearly defensible correct response?
Distractor quality	Are incorrect options plausible, parallel, and related to likely misconceptions?
Technical accuracy	Are terminology, notation, units, formulas, and supporting media correct?
Absence of clues	Do option length, wording, grammar, or lexical overlap avoid revealing the key?

Table 2: Institutional criteria used by subject-matter faculty committees during pre-operational item review.

4.2 Measures and Analysis

We calculated a validation-friction index (VFI) for each area and item source: VFI = total recorded validation comments / total submitted items. Lower values denote fewer recorded reviewer comments per item and therefore lower comment density within the workflow. VFI is used here as an administrative process indicator, not as a direct measure of item quality or validity. We also report the faculty-to-AI VFI ratio to describe relative comment density.

For the qualitative component, the project team organized recurring comments deductively using the institutional review criteria and compared patterns across knowledge areas and item sources. The available materials contained aggregate counts and area-level summaries rather than a complete item-by-comment analytic file. Accordingly, the current paper reports descriptive results without hypothesis tests, formal inter-rater agreement estimates, or item-level regression models.

5 Results

Across all areas, Gemini-drafted items received 1,810 recorded comments for 986 items (VFI = 1.84), whereas faculty-drafted items received 3,258 comments for 1,419 items (VFI = 2.30). The faculty-to-AI VFI ratio was 1.25, indicating 25.1% more recorded comments per faculty item in the aggregate. This overall comparison is descriptive and does not adjust for area, subject, cognitive level, or item format.

The direction and magnitude of the difference varied by area. In BCHS, VFI was 2.25 for Gemini drafts and 2.90 for faculty drafts; the faculty rate was 28.9% higher. In SS, the corresponding values were 1.57 and 2.16, a 37.6% higher faculty rate. HA showed the largest contrast: 1.13 for Gemini and 2.36 for faculty, or a 108.8% higher faculty rate. PMSE was the exception, with VFI of 2.01 for Gemini and 1.78 for faculty; the Gemini rate was 12.9% higher. Table 3 presents the underlying counts.

Area	AI items	AI comments	AI VFI	Faculty items	Faculty comments	Faculty VFI
PMSE	298	598	2.01	301	536	1.78
BCHS	307	691	2.25	299	866	2.90
SS	206	324	1.57	394	853	2.16
HA	175	197	1.13	425	1,003	2.36
Overall	986	1,810	1.84	1,419	3,258	2.30

Table 3: Candidate items, recorded faculty-review comments, and validation-friction index (VFI) by knowledge area and authoring source.

5.1 Recurring Qualitative Patterns

Across areas, the most consistent strength of Gemini drafts was adherence to the table of specifications. Review summaries described standardized syntax, stable metadata, and more consistent matching of expected learning outcomes and requested cognitive levels. The principal AI-related concerns were not uniform across disciplines. In PMSE, reviewers identified visual asymmetry when the key was longer or more detailed than the distractors, as well as lexical clues created by repeating words from the stem. In BCHS, some distractors were thematically distant, allowing superficial elimination rather than discrimination based on domain knowledge. In SS, comments focused on micro-format conventions and occasional elevation of cognitive complexity. In HA, most comments on Gemini items concerned fluency or stylistic refinement rather than substantive content.

Faculty-authored items showed some complementary strengths. Reviewers noted more natural disciplinary context, particularly clinical or pedagogical framing in BCHS, and stronger use of authentic misconceptions in distractor construction in some PMSE items. At the same time, faculty drafts were more heterogeneous. Recurring concerns included drift from the specified content or cognitive level, overly complex stems, implicit missing conditions, ambiguous instructions, grammatical clues, implausible distractors, and occasional incorrect or debatable keys. These patterns suggest that structured AI generation and faculty expertise address different sources of item-development risk.

5.2 Human-Platform Interface Errors

The PMSE audit highlighted a separate source of quality loss: transfer from generated text into the item-management platform. Review records traced multiple problems in exponents, subscripts, equation symbols, and graph images to manual transcription or platform entry rather than to the original Gemini output. Similar notation-transfer issues were noted in other technical subjects. This distinction matters because a workflow-level evaluation that assigns every downstream defect to the content generator will misidentify the source of error. It also suggests that parsers, direct database upload, and automated notation checks may be as important as prompt refinement.

6 Discussion

This work-in-progress study provides preliminary evidence that a persistent, blueprint-grounded Gem configuration can support large-scale item drafting with manageable expert-review burden. Gemini-drafted items had lower comment density in three of four knowledge areas and in the aggregate. The result should not be read as proof that the AI items were psychometrically superior. Rather, it suggests that explicit constraints and standardized output may reduce some forms of pre-operational variation that faculty committees otherwise need to identify and correct.

The findings are consistent with prior literature in two respects. First, LLMs can produce plausible MCQ drafts efficiently when prompts specify content and format. Second, human review remains necessary because recurring problems with keys, distractors, difficulty, and alignment do not disappear (Artsi et al., 2024; Grévisse et al., 2024; Law et al., 2025; Kaya et al., 2025). The contribution of the present case is its scale, its inclusion of four broad knowledge areas and ten subjects, its use of Mexican Spanish, and its focus on an institutional workflow rather than isolated zero-shot prompts.

The area differences are also instructive. The lowest Gemini VFI occurred in Humanities and Arts, despite the interpretive nature of those subjects. Conversely, PMSE was the only area in which Gemini drafts received more comments per item. The qualitative evidence suggests that this difference combined model-related issues (option asymmetry and lexical clues) with workflow-related notation and transfer errors. In BCHS, faculty contextual sensitivity and AI consistency were particularly complementary. These contrasts argue against a single claim that either AI or faculty authors are universally better. The appropriate question is which tasks should be standardized, automated, or reserved for expert judgment in each domain.

A defensible human-in-the-loop workflow would include at least five layers. The first is specification: the construct, blueprint cell, learning outcome, cognitive demand, item format, and intended use must be explicit. The second is controlled generation using a versioned prompt, curated source documents, and a standardized output schema. The third is automated linting for prohibited wording, option-length imbalance, key-position rules, lexical overlap, duplicate content, metadata completeness, and notation integrity. The fourth is expert review for substantive correctness, curricular meaning, bias, contextual authenticity, and distractor quality. The fifth is empirical validation after pilot or operational administration, including difficulty, discrimination, point-biserial correlations, distractor functioning, reliability, and fairness analyses. This sequence locates GenAI within an evidence-building system rather than treating generation itself as evidence of validity (American Educational Research Association et al., 2014; Kaldaras et al., 2024).

The study also suggests how VFI could mature into a more useful process measure. Future versions should separate positive confirmations from required corrections, distinguish content, cognitive-level, style, formatting, and platform errors, and record the time spent at each review stage. Item-level data would permit analyses that adjust for subject, cognitive level, format, and reviewer committee. Blinded source evaluation and matched or randomized blueprint cells would reduce expectation and selection effects. Inter-rater agreement should be estimated, and the relationship between pre-operational review indicators and post-administration item statistics should be examined.

Finally, reproducibility requires governance of the AI system itself. Institutions should retain the full system instructions, prompt revision history, knowledge-file versions, model label, generation date, original output, human edits, reviewer decisions, and platform transformations. Such provenance is necessary for auditing unexpected items and for determining whether a problem originated in the blueprint, the model, the prompt, the human editor, or the technical

interface. The project's centralized Gem architecture offers a practical starting point for this form of assessment governance.

7 Conclusion

Customized Gemini Gems produced structured MCQ drafts that, in this naturalistic institutional comparison, generated fewer recorded review comments than faculty drafts in three of four knowledge areas. Their main advantages were consistency and blueprint adherence; their recurring weaknesses involved distractor quality, verbosity, option asymmetry, and unintended clues. Faculty authors supplied disciplinary context and knowledge of authentic misconceptions, but their drafts showed greater variability in alignment and item-writing form.

The findings support neither full automation nor fully manual production. A more defensible division of labor uses AI for constrained structural drafting, automated tools for formal checks and platform transfer, and faculty committees for substantive validation and adjudication. VFI may help monitor this workflow, but only after comment types and reviewer effects are distinguished, and it cannot replace evidence from examinee responses. The next phase will test whether the observed pre-operational patterns predict difficulty, discrimination, distractor functioning, and fairness in actual administration.

Limitations

This study has several limitations. It was a retrospective process evaluation rather than a randomized comparison. AI- and faculty-authored items were not matched within blueprint cells, and source counts differed substantially in Social Sciences and Humanities and Arts. Differences in subjects, item formats, or cognitive levels may therefore partly explain the observed VFI patterns.

VFI is a novel administrative ratio with no established validity evidence. The comment count combined positive confirmations and requests for correction, may reflect reviewer communication style, and does not measure the severity of a problem. The available records did not support item-level modeling, formal qualitative reliability, inter-rater agreement, or confirmation that reviewers were blinded to item source. Actual authoring time, review time, rejection rates, and costs were not available.

No examinee response data were analyzed, so the study cannot establish equivalence in difficulty, discrimination, reliability, fairness, or score interpretations. The foundation-model version used by the Gemini application was not preserved in the available project documentation, limiting technical reproducibility. Finally, findings from one institution, one language, one prompt architecture, and one item-management workflow may not generalize to other settings or later model versions.

Ethical Considerations

The analysis reports aggregate institutional quality-assurance data and includes no student or examinee records. No operational item content, answer keys, or confidential source documents are reproduced. The prompt scaffold is abridged and translated to illustrate the workflow without compromising test security. Human committees retained final authority over item acceptance and revision. Future operational use should also include systematic review for cultural, linguistic, and construct-related bias and should preserve auditable records of AI output and human decisions.

Acknowledgments

The authors thank the faculty members who served on the subject-specific validation committees and the staff of CEIDE's educational evaluation and systems teams who supported the item-bank workflow.

References

American Educational Research Association, American Psychological Association, & National Council on Measurement in Education. (2014). *Standards for educational and psychological testing*. American Educational Research Association. https://www.testingstandards.net/uploads/7/6/6/4/76643089/standards_2014edition.pdf

Artsi, Y., Sorin, V., Konen, E., Glicksberg, B. S., Nadkarni, G., & Klang, E. (2024). Large language models for generating medical examinations: Systematic review. *BMC Medical Education, 24*, Article 354. https://doi.org/10.1186/s12909-024-05239-y

Circi, R., Hicks, J., & Sikali, E. (2023). Automatic item generation: Foundations and machine learning-based approaches for assessments. *Frontiers in Education, 8*, Article 858273. https://doi.org/10.3389/feduc.2023.858273

Gierl, M. J., Shin, J., Firoozi, T., & Lai, H. (2022). Using content coding and automatic item generation to improve test security. *Frontiers in Education, 7*, Article 853578. https://doi.org/10.3389/feduc.2022.853578

Google. (n.d.). *Tips for creating custom Gems*. Gemini Apps Help. Retrieved September 1, 2026, from https://support.google.com/gemini/answer/15235603

Grévisse, C., Pavlou, M. A. S., & Schneider, J. G. (2024). Docimological quality analysis of LLM-generated multiple choice questions in computer science and medicine. *SN Computer Science, 5*, Article 636. https://doi.org/10.1007/s42979-024-02963-6

Haladyna, T. M., Downing, S. M., & Rodriguez, M. C. (2002). A review of multiple-choice item-writing guidelines for classroom assessment. *Applied Measurement in Education, 15*(3), 309-333. https://doi.org/10.1207/S15324818AME1503_5

Haladyna, T. M., & Rodriguez, M. C. (2013). *Developing and validating test items*. Routledge. https://doi.org/10.4324/9780203850381

Kaldaras, L., Akaeze, H. O., & Reckase, M. D. (2024). Developing valid assessments in the era of generative artificial intelligence. *Frontiers in Education, 9*, Article 1399377. https://doi.org/10.3389/feduc.2024.1399377

Kaya, M., Sonmez, E., Halici, A., Yildirim, H., & Coskun, A. (2025). Comparison of AI-generated and clinician-designed multiple-choice questions in emergency medicine exam: A psychometric analysis. *BMC Medical Education, 25*, Article 949. https://doi.org/10.1186/s12909-025-07528-6

Law, A. K. K., So, J., Lui, C. T., Choi, Y. F., Cheung, K. H., Hung, K. K., & Graham, C. A. (2025). AI versus human-generated multiple-choice questions for medical education: A cohort study in a high-stakes examination. *BMC Medical Education, 25*, Article 208. https://doi.org/10.1186/s12909-025-06796-6

Linde, P., Fichter, F., Dietlein, M., Sudbrock, F., Afshar, K., Dapper, H., Fokas, E., Hillebrecht, A.-L., Raupach, T., & Laupichler, M. C. (2026). Psychometric properties and detectability of GPT-4o-generated multiple-choice questions compared with human-authored items across imaging specialties. *npj Digital Medicine, 9*, Article 132. https://doi.org/10.1038/s41746-025-02313-7

Suárez-Álvarez, J., He, Q., Guenole, N., & D'Urso, D. (2026). Using artificial intelligence in test construction: A practical guide. *Psicothema, 38*(1), 1-12. https://doi.org/10.70478/psicothema.2026.38.01

Score Comparability Limits with Common AES Validation Measures

Tim Moses
Buros Center for Testing

Seongeun Kim
National Commission on Certification of Physician Assistants

Nick Trout
Michigan State University

Abstract

This study considers measures commonly used to validate AES scoring and their limitations for indicating comparability with human rater scoring. In data simulated to reflect validation results achieved by AES national competition winners, scoring standard differences can occur across AES and human rater scoring (especially for 4-point scales vs. 2- and 3-point scales). Equipercentile methods are described and recommended for resolving the scoring differences.

1 Introduction

Automated Essay Scoring (AES) is often implemented within hybrid scoring systems, where some test takers' written responses are scored with AES and others are scored by humans (Kwako, Lottridge, & Ormerod, 2026; Palermo & Wibowo, 2024). The human scoring in hybrid systems fulfills at least two roles, to score written responses determined to be scored with difficulty or "low confidence" by AES methods, and to provide validation checks for the AES scores. To the extent that human scoring is considered the "gold standard" for the validation and interpretation of automated scoring (Palermo & Wibowo, pp. 23-24), human-machine agreement and associated measures are critical for AES validation.

What measures of human-machine agreement are most important for the validation of automated scoring? The most frequently used and reported indices include quadratic weighted kappas (QWKs, Cohen, 1968), average squared differences between scores rated by human and machine (Ramesh & Sanampudi, 2022), exact agreement percentages, and standardized mean differences (SMDs; Williamson, Xi, & Breyer, 2012). Typically, these measures are evaluated with respect to pre-defined thresholds, such as 0.70 (70%) or higher for QWKs and/or score agreements, and 0.10 or 0.15 for standardized mean differences. Although these measures address several aspects of human-machine scoring agreement, the QWK is increasingly criticized for issues like estimation sensitivity to rating scales, issues with the marginal score distributions, and inconsistencies with alternative agreement measures (Doewes, Kurdhi, & Saxena, 2023; Lewis & Casabianca, 2026). One aspect of the measures, their criticisms and proposed alternatives do not directly address is scoring comparability, or the extent to which scoring standards are comparable across the two sets of scores (i.e., human and AES).

Potential scoring standard differences could be present and not well detected in QWKs of any sets of scores (AES vs. human, human vs. human, etc.) However, these limitations are especially pronounced in AES vs. human scoring, where AES scores can exhibit reduced spreads (fewer extreme scores) and more systematic differences as compared to human vs. human disagreements that might be more random. To the extent that AES scoring is implemented in hybrid systems with no validation and adjudication procedures typically used with human scoring, any comparability issues in AES vs. human scoring that are undetected in AES validation measures like QWKs are especially consequential for entire testing populations.

The purpose of this study is to examine the limitations of common AES validation measures for ensuring comparability across AES and human scoring. For a range of rating scales (2-, 3-, and 4-point scales), we consider how different scoring standards can be for AES models that have won national competitions (Shermis, 2024). Reported

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 93–99
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

score statistics, QWKs and % exact agreements for these competition winners are reconsidered with respect to how different or incomparable the scoring standards can be. In the field of linking and equating, the primary measure for establishing comparability is when two sets of scores have equal cumulative distributions (Moses, 2025). In this study, we analyze how different the cumulative distributions of the simulated can be to reflect the human and AES score characteristics for national winners of AES scoring model competitions. We also describe and demonstrate equipercentile conversions for resolving the scoring differences shown in the simulated data as potential impacts in seemingly validated AES scoring systems.

2 Methods and Samples

The questions of interest are examined using simulated discrete distributions. We obtained the human and AES score means, standard deviations, QWKs and % exact agreement from the winners of a recent Automated Student Assessment Prize (ASAP) competition (Shermis, 2024, Tables 7.3 & 7.4) for 20 2-, 3- and 4-point essays. We simulated discrete distributions to match the scales and statistical characteristics of these essays, and also to maximize a measure of (in)comparability in the human and AES scores. The measure we considered for score incomparability is defined as an average absolute difference in cumulative proportions of human and AES scores,

$$Avg.Abs.CProporitonDifference = \sum_j p_{j,Human}\left|p_{J\leq j,Human} - p_{J\leq j,AES}\right| \quad (1)$$

where the p's indicate proportions of test takers with scores in the complete set of J scores less than or equal to score j. Equation 1 indicates the extent to which cumulative proportions of test takers' AES and human scores at each possible score differ in absolute value, averaged across the distribution of human scores. For descriptive purposes, the maximum absolute difference is also considered,

$$Max.Abs.CPercentageDifference = 100\{Max\left|p_{J\leq j,Human} - p_{J\leq j,AES}\right|\} \quad (2)$$

To the extent that scoring standards differ (e.g., higher scores are more difficult to achieve with AES vs. humans and exhibit larger cumulative proportions), values for Equations 1 and 2 will be larger and more indicative of different or incomparable scoring across AES and human systems. By generating bivariate distributions that exhibited the agreements achieved in national AES scoring competitions and also maximize Equation 1, we could consider how incomparable scoring standards could get when AES scoring that meets validation thresholds and wins national competitions.

To the extent that AES scores and human rater scores are systematically different and reflect differences in leniency or stringency, the raw AES and human scores should have at least some score points with different cumulative percentages. These differences in cumulative percentages indicate that the same score point may represent a different relative position in the AES and human score distributions. As mentioned, two sets of scores are considered comparable when they have equal cumulative distributions in linking and equating (Moses, 2025). Therefore, to address differences in the cumulative distributions of AES and human rater scores, equipercentile conversions were produced. Specifically, equipercentile conversions of AES scores to the human score scales are computed by identifying scores on the human scale with the same cumulative percentages as the corresponding AES scores (approximately "equal" percentiles).

3 Results

Table 1 presents overall results, including the human and AES score statistics and agreement measures from Shermis (2024), and the measures of scoring incomparability (Equations 1 and 2) Results show that for national competition winners with QWKs almost always above established thresholds (>0.70) and with relatively small differences in mean scores and SMDs, consequential differences in scoring standards can occur. The potential for scoring standard differences is up to 3% for essays with 2-point scales, 6% for essays with 3-point scales, and 7% for essays with 4-point scales. We consider both proportions and cumulative proportions of scores of the score distributions in detail for all essay items, depicted in Figures 1-3.

Table 1. *Human and AES score statistics and agreement measures for national AES competition winners from Shermis (2024), and Avg.Abs.CProportionDifferences and Max.Abs.CPercentageDifferences.*

	From Shermis (2024)							*Current Study*	
Item	**Score Range**	**Mean (Human)**	**SD (Human)**	**Mean (AES)**	**SD (AES)**	**%Exact Agreement**	**QWK**	***Avg.Abs. CProporitonDifference***	***Max.Abs. CPercentageDifference***
DR04_1715RE1T10_05	1-2	1.23	0.42	1.24	0.43	0.95	0.88	0.01	1%
DR04_1715RE4T05G04_07	1-2	1.45	0.50	1.46	0.50	0.94	0.87	0.01	1%
DR08_1715RE2T13_05	1-2	1.66	0.48	1.69	0.46	0.91	0.79	0.01	3%
DR08_1715RE4T05G08_07	1-2	1.80	0.40	1.81	0.39	0.94	0.80	0.00	2%
DR04_1715RE4T05G04_03	1-3	1.55	0.73	1.57	0.74	0.89	0.90	0.01	2%
DR04_1715RE4T05G04_09	1-3	1.55	0.67	1.55	0.67	0.78	0.75	0.01	1%
DR04_1715RE4T08G04_03	1-3	1.96	0.71	1.95	0.70	0.91	0.89	0.01	2%
DR04_1715RE4T08G04_06	1-3	1.56	0.80	1.57	0.80	0.96	0.97	0.01	1%
DR04_1715RE4T08G04_09	1-3	1.46	0.57	1.48	0.56	0.87	0.80	0.02	3%
DR08_1715RE2T13_08	1-3	1.99	0.73	2.02	0.73	0.75	0.69	0.01	2%
DR08_1715RE4T05G08_03	1-3	2.22	0.88	2.22	0.89	0.90	0.92	0.01	1%
DR08_1715RE4T05G08_09	1-3	2.20	0.75	2.15	0.72	0.81	0.81	0.02	6%
DR08_1715RE4T08G08_03	1-3	2.46	0.69	2.44	0.68	0.92	0.89	0.01	2%
DR08_1715RE4T08G08_06	1-3	2.42	0.85	2.41	0.85	0.96	0.97	0.00	1%
DR08_1715RE4T08G08_09	1-3	1.94	0.66	1.97	0.63	0.81	0.77	0.01	4%
DR04_1715RE4T05G04_06	1-4	1.90	1.05	1.90	1.06	0.82	0.90	0.03	7%
DR04_1715RE4T08G04_07	1-4	2.02	0.99	2.01	0.99	0.94	0.97	0.02	4%
DR08_1715RE2T13_07	1-4	1.65	0.84	1.68	0.91	0.83	0.89	0.03	5%
DR08_1715RE4T05G08_06	1-4	2.86	1.11	2.90	1.10	0.77	0.87	0.01	2%
DR08_1715RE4T08G08_07	1-4	3.02	0.95	3.03	0.95	0.93	0.96	0.01	3%

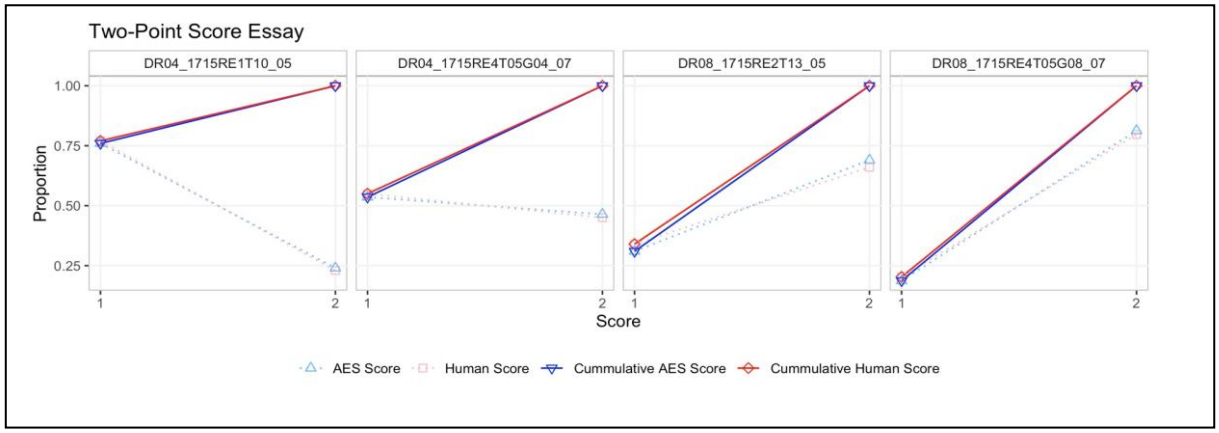


Figure 1. AES and Human Score Distributions (Essays with 2-point scales).

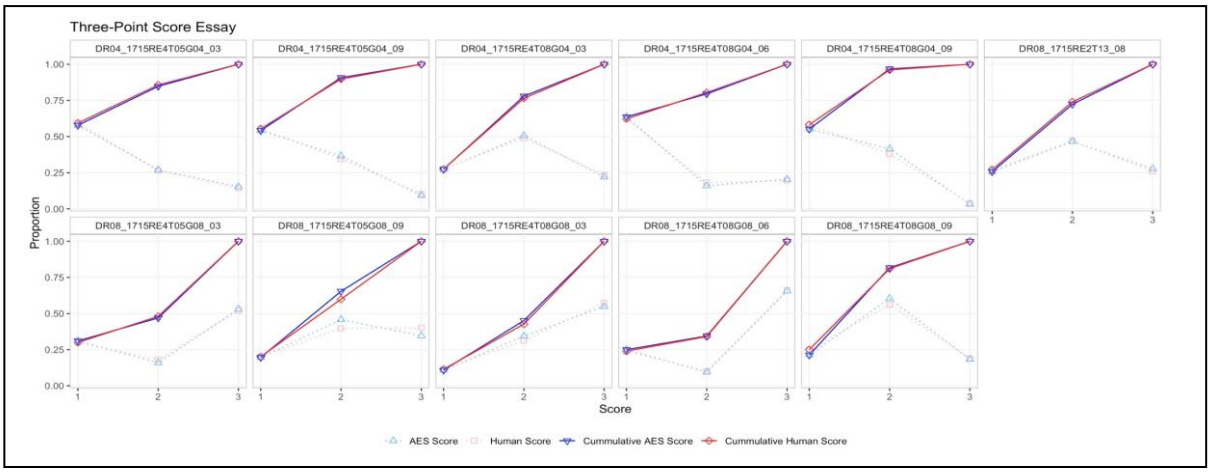


Figure 2. AES and Human Score Distributions (Essays with 3-point scales).

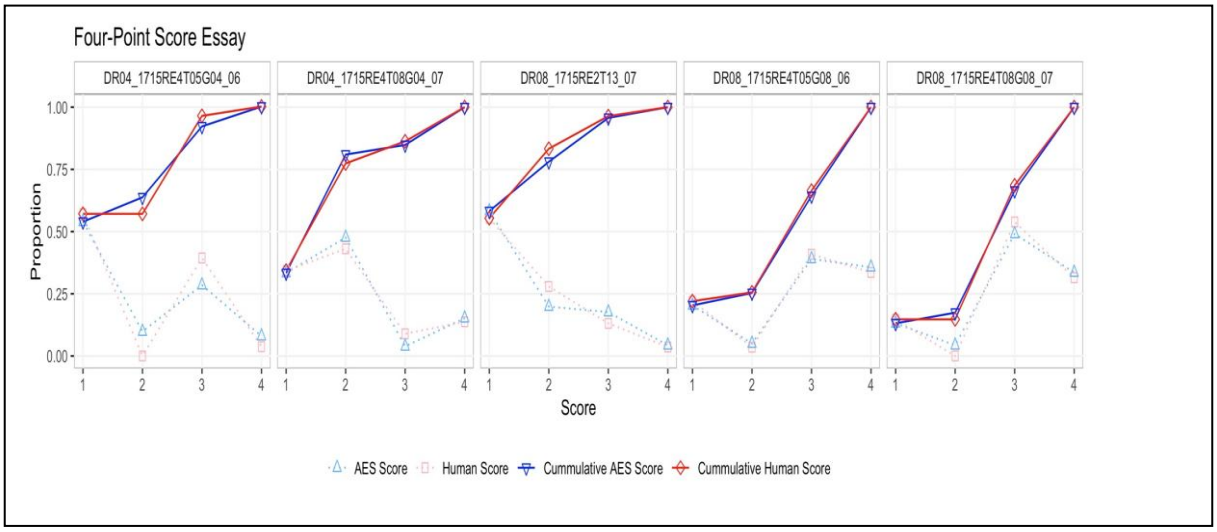


Figure 3. AES and Human Score Distributions (Essays with 4-point scales).

Overall, the proportions and cumulative proportions of AES and human scores in Figures 1-3 show two patterns. First, differences between AES and human score distributions were more apparent as the number of score points increased. Proportions and cumulative proportions were more similar for 2- and 3-point essays, and exhibited greater differences for 4-point essays. Second, human ratings tended to be more stringent than AES ratings, although the pattern varied across score categories. For the 2-point essays, human raters assigned a greater proportion of scores to 1 than AES, whereas AES assigned a greater proportion to score 2. For the 3point essays, AES assigned a greater proportion of scores to higher score categories than human raters. For the 4-point essays, the differences between human and AES ratings varied across score categories. Between the two extreme score categories (1 and 4), AES assigned a greater proportion of scores to score 4, whereas between the two middle score categories, AES assigned a greater proportion of scores to score 2.

As previously mentioned, two sets of scores are considered comparable when their cumulative distributions are equal (Moses, 2025). Therefore, AES scores were linked to human score scales, reflecting a treatment of human scores as the gold standards. These linkings were computed as unsmoothed equipercenitle conversions. The equipercentile results are summarized in Figures 4-6, which show the differences in the AES-to-Human equipercentile converted scores and the Human scores across the essay score scales. Consistent with the differences observed in the AES and human score distributions (Figures 1-3), the average absolute difference between AES-to-Human converted scores and Human scores increased as the number of score points increased: from 0.022 for the 2-point essays, to 0.029 and 0.135 for the 3- and 4-point essays. As expected, when AES scores were converted to the scales of the human raters, the converted AES scores tended to be lower than the raw and unconverted AES scores. The differences were relatively small for the 2- and 3-point essays but were more apparent for the 4-point essays. In particular, differences greater than 0.5 score points were observed for the first (DR04_1715RE4T05G04_06) and last (DR08_1715RE4T08G08_07) 4-point essays. For these items, the proportions of responses assigned a score of 2 by human raters were close to 0%, which likely contributed to the relatively large converted score differences.

4 Conclusions

This study contributes to discussions about limitations of commonly used validation measures for AES (Doewes, Kurdhi, & Saxena, 2023; Lewis & Casabianca, 2026). For AES systems that meet thresholds for agreement and even win national score model competitions, meaningful differences in AES vs. human scoring standards can occur. These score comparability issues have implications for AES uses in the field, including hybrid scoring systems for state testing where AES scoring can be shown to meet established validation thresholds with respect to human scoring but still exhibit comparability issues like few or no highest possible scores, different score variances, or other scoring discrepancies.

These findings highlight the need to consider the magnitude of scoring (in)comparability directly, such as through distributions and equipercentile conversions, even when agreements exceed acceptable thresholds for AES validation measures (e.g., QWKs, % exact agreements, and SMDs). As illustrated in Figures 1-6, even when QWKs were high enough to win AES competitions, differences could be observed in the distributions of AES and human rating scores (Figures 1-3), and in equipercentile score conversions (Figures 4-6). These discrepancies tended to be larger as the number of score points increased, especially for essays with 4-point scales. The larger differences could be practically meaningful and affect score interpretations. Therefore, it is important to not rely solely on QWKs but also to consider the magnitude of scoring differences at each score point. A final implication of our results is that unrounded AES and human score scales may be useful to consider for hybrid scoring systems. Scoring differences of about 0.2 score points were found in some of our equipercentile results. Although these differences may not be observed in scores reported after rounding, they are large enough that discrepant rounded score distributions could occur. Unrounded equipercentile converted scores may be especially useful for the reporting of essay scores in hybrid scoring systems, and for addressing AES vs. human scoring differences not indicated in frequently used AES validation measures.

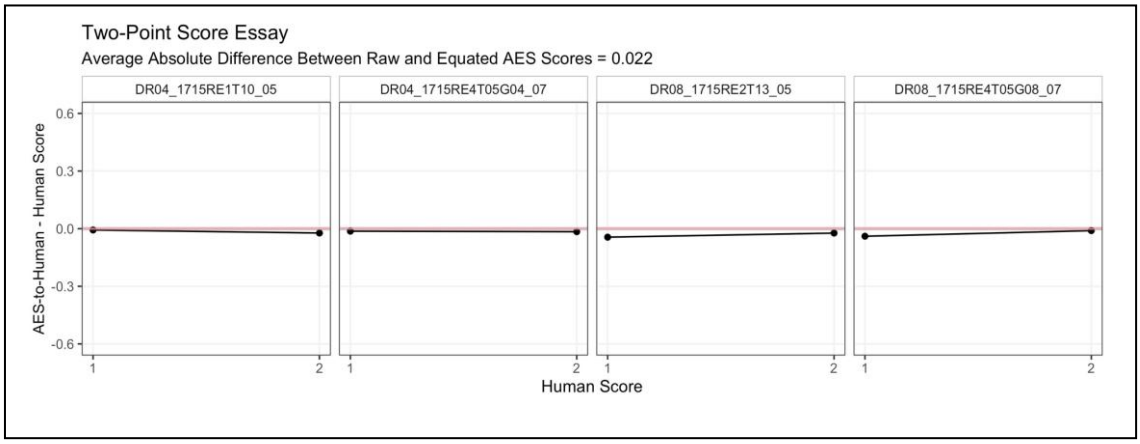


Figure 4. Differences in AES-to-Human Equipercentile and Human Scores (Essays with 2-point scales).

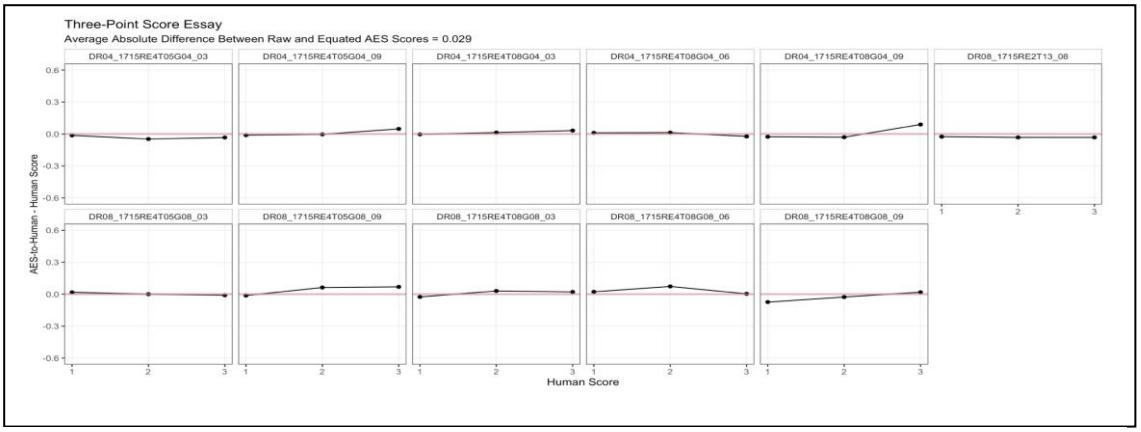


Figure 5. Differences in AES-to-Human Equipercentile and Human Scores (Essays with 3-point scales).

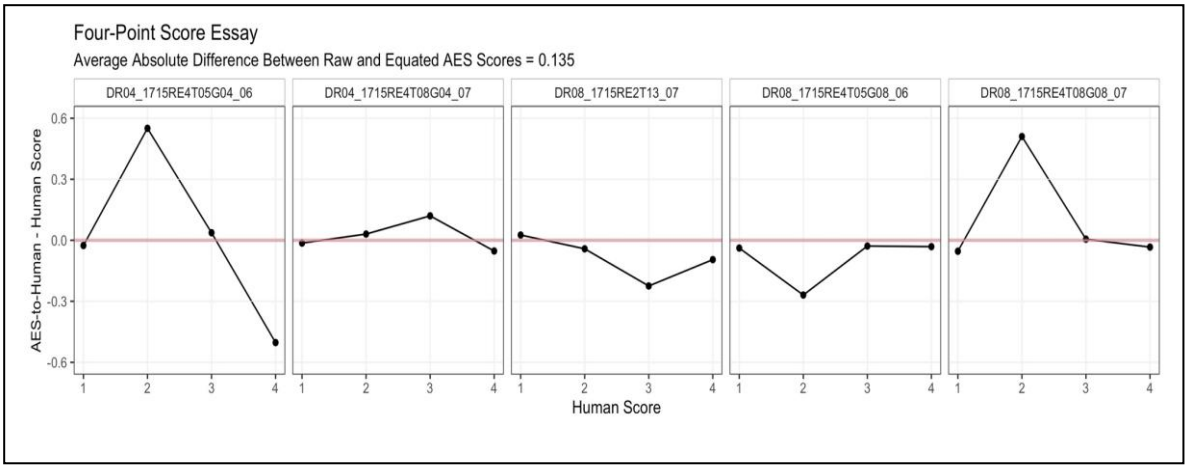


Figure 6. Differences in AES-to-Human Equipercentile and Human Scores (Essays with 4-point scales).

References

Cohen, J. (1968). Weighted kappa: Nominal scale agreement provision for scaled disagreement or partial credit. *Psychological Bulletin, 70(4),* 213-220. https://doi.org/10.1037/h0026256

Chung, K. W. K. & O'Neil, H. F. (1997). *Methodological approaches to online scoring of essays* (ERIC reproduction service no ED 418 101).

Doewes, A., Kurdhi, N. A., & Saxena, A. (2023). Evaluating quadrative weighted kappa as the standard performance metric for automated essay scoring. In M. Feng, T. Kaser, & P. Talukdar (Eds.), *Proceedings of the 16th International Conference on Educational Data Mining* (pp. 101-113). International Educational Data Mining Society (IEDMS). https://doi.org/10.5281/zenodo.8115784

Huang, Y., Palermo, C., & Wilson, J. (2026). Accuracy and fairness of generative AI in automated essay scoring: Comparing GPT-4o, feature-based models, and human raters. *Assessing Writing, 69,* 101047. https://doi.org/10.1016/j.asw.2026.101047

Kwako, A., Lottridge, S., & Ormerod, C. (2026). Using confidence modeling to optimize overall score quality in hybrid scoring systems. *Educational Measurement: Issues and Practice, 45(2),* https://doi.org/10.1111/emip.70019

Lewis, J., & Casabianca, J. M. (2026). Limitations of QWK in evaluating automated and human scoring systems. *Educational Measurement: Issues and Practice, 45(1),* https://doi.org/10.1111/emip.70017

Moses, T. (2025). Scaling, equating, and linking. In L. L. Cook & M. J. Pitoniak (Eds.), *Educational measurement* (5th ed., pp. 735-820). National Council on Measurement in Education. https://doi.org/10.1093/oso/9780197654965.002.0003

Palermo, C., & Wibowo, A. (2024). Automated essay evaluation at scale: Hybrid automated scoring/hand scoring in the summative assessment program. In Shermis, M. D., & Wilson, J. (Eds.). *The Routledge International Handbook of Automated Essay Evaluation* (pp. 23-39). Routledge. https://doi.org/10.4324/9781003397618-3

Ramesh, D., & Sanampudi, S. K. (2022). An automated essay scoring systems: a systematic literature review. *Artificial intelligence review, 55(3),* 2495-2527. https://doi.org/10.1007/s10462-021-10068-2

Ramineni, C., & Williamson, D. M. (2013). Automated essay scoring, Psychometric guidelines and practices. *Assessing Writing, 18(1),* 25-39.

Shermis, M. D. (2024). Automated scoring for NAEP short-form constructed responses in reading. In Shermis, M. D., & Wilson, J. (Eds.). *The Routledge International Handbook of Automated Essay Evaluation* (pp. 117-140). Routledge. https://doi.org/10.4324/9781003397618-9

Williamson, D. M., Xi, X., & Breyer, F. J. (2012). A framework for evaluation and use of automated scoring. *Educational Measurement: Issues and Practice, 31(1),* 2-13. https://doi.org/10.1111/j.1745-3992.2011.00223.x

Finding Evidence of Communication Behaviors using LLMs

Tazin Afrin[1], Saed Rezayi[1], Janet Mee[1], Polina Harik[1], Le An Ha[2]

[1]National Board of Medical Examiners, Philadelphia, USA
{tafrin, srezayidemne, jmee, pharik}@nbme.org
[2]Ho Chi Minh City University of Foreign Languages, Vietnam
anhl@huflit.edu.vn

Abstract

This study compares two automated scoring approaches to identify communication behaviors in physician responses to patient questions: prompt-based scoring with Large Language Models (LLM) and supervised transformer-based models. Our results show that although transformer-based models provide more consistent performance, competitive LLM performance is promising in reducing the need for extensive annotated training data.

1 Background

Automated scoring of text-based responses has been used to improve the efficiency, scalability, and reliability of assessments for a variety of purposes (Shermis and Burstein, 2013). Historically, operational scoring systems have relied heavily on supervised machine learning models trained on human-scored responses. Encoder-based transformer models such as BERT, RoBERTa, and DeBERTa have demonstrated strong performance across scoring and classification tasks by leveraging large-scale pretraining followed by fine-tuning on human-annotated assessment data (Devlin et al., 2019; Liu et al., 2019; He et al., 2020; Ormerod et al., 2021; Suen et al., 2023; Rezayi et al., 2025). However, building these systems usually requires significant investment in annotation, reconciliation, and training (Bejar et al., 2016; Kaldaras and Haudek, 2022).

Generative large language models (LLMs) have introduced a new alternative to automated scoring. Models such as GPT, Claude, Gemini and Llama are trained primarily through autoregressive next-token prediction on massive text corpora rather than through task-specific supervised training for scoring or classification (Brown et al., 2020; Gilardi et al., 2023). Unlike encoder-based transformers, generative LLMs learn broad linguistic and contextual knowledge during pretraining and can often be adapted to new tasks through prompting alone. Recent research suggests that these models can perform complex categorization and scoring tasks with little or no task-specific training data (Wei et al., 2022; Kojima et al., 2022; Gilardi et al., 2023). Their effectiveness has been further enhanced through instruction tuning, reinforcement learning from human feedback (RLHF), and prompting strategies, encouraging explicit reasoning and structured decision-making (Ouyang et al., 2022; Wei et al., 2022).

The present study compares two automated scoring paradigms for identifying evidence of patient-centered communication behaviors (PCBs) in transcribed physician responses to patient questions. For example, a span of text such as "it is quite understandable that you are frustrated by this situation" would be extracted as evidence of a PCB reflecting acknowledgment of a patient's emotions. We examine prompt-based scoring using generative large language models (LLMs), which unlike a supervised transformer-based model does not rely on substantial annotated training data and task-specific model fitting. Beyond comparing scoring performance between the two types of models, the study explores the implications of these approaches for the future development of automated scoring systems.

2 Aims

This study addresses the following research questions:

- How do prompt-based LLMs compare with fine-tuned encoder-based transformer model for PCB span extraction?
- Which modeling approach demonstrates the greatest potential for scalable automated assessment?

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 100–103
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

3 Sample(s)

The dataset consists of transcribed spoken responses produced by medical students interacting with an online communication learning environment. Students viewed video vignettes of the patient scenario and responded verbally to the prompt "What would you say next?" The student responses were then automatically transcribed.

Each student response can contain multiple examples of patient-centered communication behaviors (PCBs). PCBs are defined through functional communication categories, learning objectives, and behavioral definitions such as acknowledging emotion, explaining treatment, or clarifying instructions. Human-annotated training and evaluation examples included character-level span annotations indicating the portion of the response representing one PCB. The present analysis includes 14 cases representing a range of clinical scenarios. Each case contains between 2 and 6 PCBs and represents a distinct PCB detection scenario with varying contextual and linguistic complexity.

4 Methods

The task was formulated as a span extraction problem. Given a student response transcript and a target PCB definition, models were required to:

- Determine whether the PCB was present in the response.
- Return character-level indices corresponding to the PCB text span.

It is possible that no PCB was identified in that response.

4.1 Models

For the supervised model, a DeBERTa-based transformer model was fine-tuned for PCB span extraction using annotated examples (Rezayi et al., 2025). The model leveraged contextualized token representations to predict PCB-relevant spans within student responses.

For the LLM based approach, GPT-5 was evaluated using a structured few-shot prompting framework. The model was instructed to identify whether the PCB was fulfilled explicitly or implicitly in the student response, and return the extracted text span with character-level indices.

We also used Claude-Sonnet-4.6 using the same prompting framework and examples as GPT-5 to enable direct comparison across LLMs.

The **prompting framework** for both of the LLM models emphasized:

- Contextual grounding using clinical scenarios
- Explicit PCB definitions and examples
- Demonstrations of valid and invalid PCB matches
- Student response to evaluate
- Output formatting requirements

Two **evaluation metrics** were used:

Character-level F1 overlap: Measures span overlap F1-score between predicted and gold annotations at the character level.

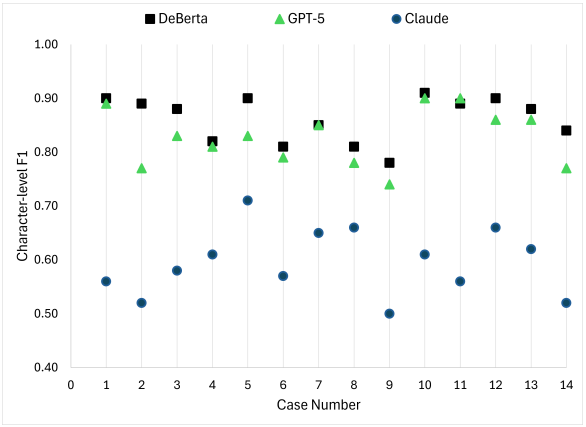
Binary F1-score: Measures whether the system correctly identified the presence or absence of the PCB regardless of exact span boundaries.

5 Results

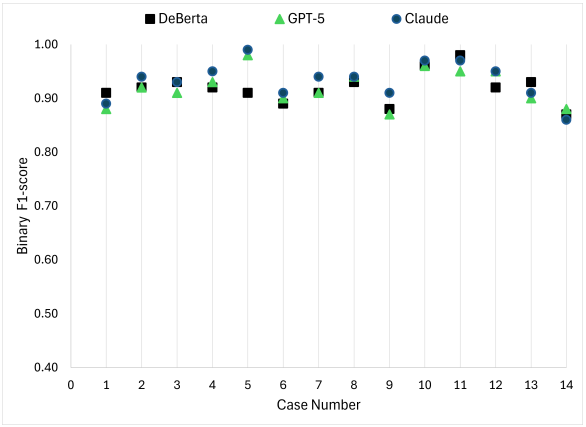
Figure 1 summarizes model performance across the 14 communication cases. The DeBERTa-based model consistently achieved the strongest character-level overlap performance across most cases, with scores generally ranging from 0.81 to 0.91 with an average of 0.86. GPT-5 achieved competitive but slightly lower span overlap performance with average 0.83, while Claude-Sonnet demonstrated substantially lower character-level overlap scores, average 0.60, despite strong binary detection performance. In contrast, binary PCB detection performance was high across all three systems. Claude-Sonnet frequently achieved the highest binary F1 scores despite weaker span localization. These findings suggest that supervised fine-tuning supports precise localization of PCB spans.

GPT-5 demonstrated competitive binary and span-level performance without task-specific fine-tuning approaching DeBERTa performance in several cases. The model appeared more variable across cases, particularly for precise span localization. These results suggest that prompt-based LLMs can approximate supervised span extraction performance under carefully designed prompting conditions.

The findings indicate a divergence between PCB identification and precise span extraction. Encoder-based supervised models appear better suited for fine-grained span localization, whereas frontier LLMs demonstrate strong semantic recognition



(a) Character-level F1 scores



(b) Binary F1 scores

Figure 1: Performance comparison across models.

abilities without fine-tuning. The results also suggest that evaluation metric choice substantially affects conclusions about model quality.

6 Conclusions

This study demonstrates the feasibility of automated patient-centered communication assessment using both supervised transformer models and generative LLMs. The DeBERTa-based supervised model currently provides the strongest and most consistent span localization performance, supporting value of task-specific fine-tuning for precise PCB extraction. On the other hand, GPT-5 achieved competitive performance using only few-shot prompting, suggesting that modern LLMs may reduce the need for extensive supervised training data in communication assessment tasks. Claude-Sonnet's stronger binary detection but weaker span finding highlights the importance of distinguishing between semantic recognition vs span extraction performance.

Future work will include larger-scale evaluation across individual PCB categories, error analysis of binary versus span overlap, human-model agreement analysis, and finally prompt optimization studies. Overall, this work contributes to ongoing efforts to develop scalable, automated assessment systems for physician-patient communication training in medical education.

References

Isaac I Bejar, Robert J Mislevy, and Mo Zhang. 2016. Automated scoring with validity in mind. *The Wiley handbook of cognition and assessment: Frameworks, methodologies, and applications*, pages 226–246.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, and 1 others. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, pages 4171–4186.

Fabrizio Gilardi, Meysam Alizadeh, and Maël Kubli. 2023. Chatgpt outperforms crowd workers for text-annotation tasks. *Proceedings of the National Academy of Sciences*, 120(30):e2305016120.

Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2020. Deberta: Decoding-enhanced bert with disentangled attention. *arXiv preprint arXiv:2006.03654*.

Leonora Kaldaras and Kevin C Haudek. 2022. Validation of automated scoring for learning progression-aligned next generation science standards performance assessments. In *Frontiers in education*, volume 7, page 968289. Frontiers Media SA.

Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. In *Advances in Neural Information Processing Systems*, volume 35, pages 22199–22213.

Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.

Christopher M Ormerod, Akanksha Malhotra, and Amir Jafari. 2021. Automated essay scoring using efficient

transformer-based language models. *arXiv preprint arXiv:2102.13136*.

Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. Training language models to follow instructions with human feedback. In *Advances in Neural Information Processing Systems*, volume 35, pages 27730–27744.

Saed Rezayi, LA Ha, Yiyun Zhou, Andrew Houriet, Angelo D'Addario, Peter Baldwin, Polina Harik, and Victoria Yaneva. 2025. Automated scoring of communication skills in physician–patient interaction: Balancing performance and scalability. In *Proc 20th BEA Workshop*.

Mark D Shermis and Jill Burstein. 2013. Handbook of automated essay evaluation. *NY: Routledge*.

King Yiu Suen, Victoria Yaneva, Janet Mee, Yiyun Zhou, Polina Harik, and 1 others. 2023. Acta: Short-answer grading in high-stakes medical exams. In *Proceedings of the 18th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2023)*, pages 443–447.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc V. Le, and Denny Zhou. 2022. Chain-of-thought prompting elicits reasoning in large language models. In *Advances in Neural Information Processing Systems*, volume 35, pages 24824–24837.

Designing and Evaluating a Measurement-Oriented RAG Framework for ELL Writing Assessment

Daeryong Seo
Pearson*
daeryong.seo@pearson.com

Abstract

This design-focused study compared measurement-oriented ELL_RAG with conventional RAG for writing-assessment analytics. ELL_RAG showed a descriptive advantage overall and on global/population-level tasks, while local-evidence performance was nearly equivalent. Findings support task–evidence alignment rather than universal superiority and highlight tool orchestration as a continuing design challenge.

1 Introduction

Retrieval-augmented generation (RAG) grounds language-model output in external evidence and has proved useful when relevant information can be retrieved directly from a knowledge source (Lewis et al., 2020). More recent RAG architectures have become increasingly modular and domain-specific (Gao et al., 2023). These developments are relevant to educational writing assessment, where large language models can support scoring, feedback, and interpretation but must do more than generate fluent explanations. Assessment claims require evidence appropriate to their intended interpretations and uses, including attention to validity, reliability, fairness, and uncertainty (Messick, 1989; Kane, 2013; Pack et al., 2024). A response may be fluent and plausible yet still be weak if the evidence used is not appropriate for the conclusion. For example, retrieved individual essays cannot by themselves support population-level claims.

This creates an important evidence-design problem. Some questions are local and can be answered from an official rubric, structured metadata, or a retrieved essay. Others require full-dataset statistics, inter-rater reliability, controlled subgroup comparisons, or other population-level analyses. A semantically relevant essay may therefore be useful context while still being inappropriate evidence for a dataset-level claim. Conversely, deterministic statistics do not replace the interpretive value of a local response when the question concerns a specific essay or rubric application. The broader RAG literature reflects a similar distinction: Edge et al. (2024) contrasted local retrieval with global corpus-level sensemaking, and Murtaza et al. (2025) combined document retrieval with database function calls for structured-data questions.

The present study addresses this problem by designing ELL_RAG, a research prototype for ELL writing-assessment analytics. The architecture combines deterministic measurement tools, constrained essay retrieval, rubric lookup, and LLM orchestration under a hierarchical evidence design: dataset-level claims come from deterministic analyses, official score meanings from rubric evidence, and retrieved essays from local contextual evidence. ELL_RAG is compared with a conventional RAG baseline to examine whether these additional mechanisms provide incremental value when population-level evidence is required while preserving performance when retrieval alone may be sufficient. The emphasis is therefore architectural design and task–evidence alignment rather than a general contest between two systems.

1.1 Study Purpose and Research Questions

The primary aim was to design and examine a measurement-oriented RAG workflow that selects retrieval, rubric, deterministic measurement, or combined evidence pathways according to the level of inference requested (Appendix A). The comparison with Conv_RAG was used

* This study describes the author's research prototype developed using an open-source dataset. It is not an operational Pearson scoring system or client implementation and is not intended to replace human judgment, psychometric review, validation, or established governance processes.

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 104–112
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

diagnostically to identify where conventional retrieval remained sufficient, where measurement-oriented mechanisms added support, and what strengths or limitations characterized each architecture. The evaluation was intentionally descriptive and served as evidence about whether the proposed design behaved as intended.

Accordingly, the study addressed three research questions:

1. How does ELL_RAG differ from Conv_RAG in supporting evidence selection, integration, and interpretation across the 15 prespecified questions compared?
2. How do ELL_RAG and Conv_RAG differ in their performance on local, directly retrievable evidence tasks versus global/population-level measurement tasks?
3. How do ELL_RAG and Conv_RAG differ in their strengths and limitations related to evidence selection, measurement interpretation, uncertainty control, scope alignment, and tool orchestration?

2 Method

2.1 Sample(s)

The study used 5,840 essays from the ELLIPSE Corpus, which contains approximately 6,500 ELL essays from U.S. statewide testing across 44 prompts. Essays were scored by two trained raters on Overall proficiency and six analytic dimensions, with accompanying student demographic variables (Crossley et al., 2023). The analytic sample included 5,840 complete cases used in the RAG analyses.

2.2 Study Design

The study used one frozen, design-focused parallel comparison of Conv_RAG and ELL_RAG. Both systems independently answered the same 15 prespecified questions without conversational history, producing 30 responses. The frozen protocol prevented iterative adjustment of questions or prompts after observing system outputs. The study emphasized architecture, evidence-pathway behavior, and preliminary response quality rather than inferential claims of universal superiority (Appendix A).

2.3 Systems Compared

Conv_RAG used conventional semantic retrieval and generation over corpus documents, rubric material, structured metadata, and essays. It did not include dedicated deterministic functions for full-dataset summaries, reliability estimation, or controlled subgroup analyses. ELL_RAG retained local retrieval but added measurement-aware routing, deterministic analytical tools, authoritative rubric lookup, structured eligibility controls, and an explicit local-versus-global evidence hierarchy. Dataset-level questions could therefore be routed to measurement functions, while local questions could continue to use retrieval, rubric, or metadata pathways.

2.4 Matched Design Conditions

The systems used the same ELLIPSE corpus of 5,840 essays (Crossley et al., 2023), sentence-transformers/all-MiniLM-L6-v2 embeddings, GPT-4o base model, generation settings, final evidence depth, and frozen question set. The intended difference was architectural: Conv_RAG used conventional retrieval, whereas ELL_RAG added measurement-support pathways and a larger initial retrieval pool before filtering/reranking to the final evidence set.

2.5 Evaluation Question Set

The 15 questions were developed based on a review of relevant assessment, measurement, and retrieval-augmented analytics literature and standards and were distributed evenly across five assessment-oriented categories spanning 1) validity, fairness, and quality; 2) subgroup analysis and oversight; 3) measurement and retrieval evidence; 4) score interpretation; and 5) evidence and uncertainty. The question-development framework drew on established guidance for assessment validity, reliability, fairness, and score interpretation, the ELLIPSE assessment context, and literature concerning the integration of retrieved and structured evidence (AERA, APA & NCME, 2014; Cohen, 1968; Crossley et al., 2023). Five items (Q02, Q04, Q05, Q07, and Q13) were prespecified global/population-level probes requiring dataset-level evidence or integration with such evidence; three (Q08, Q09, and Q11) were local-evidence probes centered on appropriate use of essays, rubrics, or metadata. These were descriptive task groupings, not psychometric subscales (Appendices B and C).

2.6 Scoring Framework

The study author scored all 30 responses on four criteria, each from 0 to 2: Evidence Use, Grounding, Accuracy, and Scope Alignment (C1); Measurement Interpretation (C2); Uncertainty and Overclaim Prevention (C3); and Clarity and Completeness (C4). The maximum response score was 8. Scores of 0, 1, and 2 represented major failure, partial success with a material limitation, and accurate/appropriately bounded performance, respectively.

For C1, for example, a score of 0 indicated a fundamental evidence-source or scope failure, such as substituting local essays for a required population-level analysis; 1 indicated meaningful but incomplete or materially flawed evidence use; and 2 required accurate, appropriately sourced, scope-aligned evidence. Therefore, C1–C4 constituted the complete formal scoring framework.

2.7 Analysis

Analysis was descriptive and design-focused. System totals, percentages, mean item scores, C1–C4 means, category summaries, and item differences were examined. Global/population-level and local-evidence probes were summarized separately to test the study's central design expectation that architecture behavior would depend on evidentiary demand. Scoring rationales were also reviewed for recurring strengths and limitations such as safe abstention, local-to-global scope mismatch, routing error, and incomplete multi-source integration. Because the study used only 15 questions and one author-rater, no inferential test of superiority was used, and category means were treated only as summaries of the sampled items.

3 Results

Results are interpreted as preliminary evidence about architectural behavior under different evidentiary demands.

3.1 RQ1. Overall Response Quality and Performance Across C1–C4

Across 15 items, ELL_RAG scored 105/120 (87.5%; mean = 7.00/8), compared with 88/120 (73.3%; mean = 5.87/8) for Conv_RAG. This represents a descriptive advantage for ELL_RAG of 17 total points, 14.2 percentage points, or 1.13 points per item. ELL_RAG received eight perfect item scores and Conv_RAG six (Table 1). These totals summarize the frozen task set but do not by themselves establish general system superiority.

System / Difference	Total Score	Percent	Mean Item Score	8/8 Items
Conv_RAG	88/120	73.3%	5.87/8	6
ELL_RAG	105/120	87.5%	7.00/8	8
ELL_RAG − Conv_RAG	+17	+14.2 pp	+1.13	+2

Table 1: Overall Descriptive Performance Across the 15 Items

The largest criterion difference occurred for C1 (1.60 vs. 1.13), followed by C2 (1.80 vs. 1.53) and C3 (1.93 vs. 1.67); C4 was closest (1.67 vs. 1.53). Thus, the main descriptive contrast concerned evidence use, scope, and measurement interpretation more than response clarity (Figure 1).

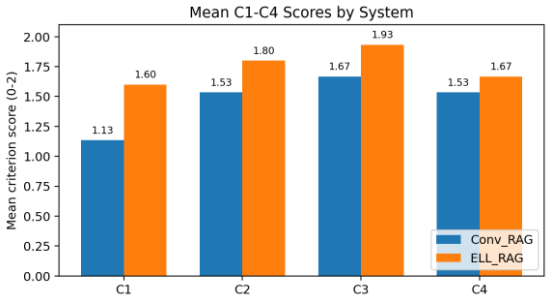


Figure 1: Mean C1–C4 scores by system. Each criterion was scored from 0 to 2.

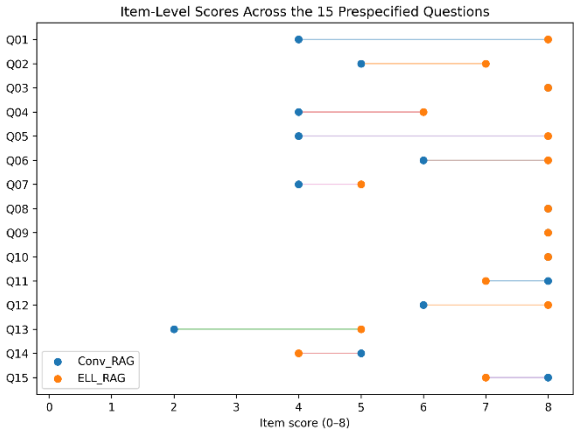


Figure 2: Item-level response scores across Q01–Q15.

ELL_RAG scored higher on eight questions, Conv_RAG on three, and four were ties (Figure 2). The largest category differences favored ELL_RAG in Subgroup fairness/transparency/human oversight (7.33 vs. 4.67) and Validity/fairness/quality (7.67 vs. 5.67); differences in the remaining three categories were only 0.33 points (Table 2). Because each category

contained three items, these values are descriptive rather than stable category effects.

Evaluation Category	Conv_RAG	ELL_RAG	Difference
Validity, fairness, and quality in assessment interpretation	5.67	7.67	+2.00
Subgroup fairness, transparency, and human oversight	4.67	7.33	+2.67
Integration of measurement and retrieval evidence	6.67	7.00	+0.33
Responsible score interpretation and professional use	7.33	7.67	+0.33
Ethical handling of evidence, uncertainty, and overclaiming	5.00	5.33	+0.33

Table 2: Category-Level Mean Item Scores

Note. Scores are means on the 0–8 item scale. Difference = ELL_RAG minus Conv_RAG. Each category contained three items; category means are therefore descriptive summaries of the sampled tasks rather than estimates of separate subscales.

3.2 RQ2. Performance Under Global/Population-Level Versus Local Evidentiary Demands

3.2.1 Global and Population-Level Evidence

On the five global/population-level probes, ELL_RAG obtained 31/40 (mean = 6.20/8) versus 19/40 (3.80/8) for Conv_RAG. C1 means were 1.20 and 0.20, respectively. These questions required full-dataset evidence, reliability information, controlled subgroup comparisons, or explicit integration of dataset-level evidence with other sources. The difference is therefore best interpreted as performance under a specific evidentiary demand rather than as a general ranking.

Conv_RAG often recognized when retrieved context could not support the requested statistic and declined to fabricate unavailable population-level results, demonstrating useful safe abstention. In several cases it could describe what analysis would be needed but could not provide the requested quantitative evidence. ELL_RAG more often supplied such evidence because it could route the question to deterministic measurement functions. This distinction reflects architecture rather than simple retrieval quality.

ELL_RAG nevertheless showed orchestration weaknesses. Q07 used an unnecessary analytical framing, Q13 omitted requested numerical results, and Q14 incompletely integrated multiple evidence sources. These cases indicate that tool availability did not always translate into correct routing, evidence integration, or final interpretation.

3.2.2 Local Evidence

The local-evidence probes showed near-equivalent performance: Conv_RAG scored 24/24 and ELL_RAG 23/24, with both averaging 2.00/2 on C1 and C3. Both appropriately prioritized the official Vocabulary rubric over a conflicting retrieved example on Q08 and used structured prompt metadata to determine eligibility on Q09. On Q11, Conv_RAG was one point higher because it more explicitly stated that an authorized human review process, rather than the system alone, would be required before changing an official score.

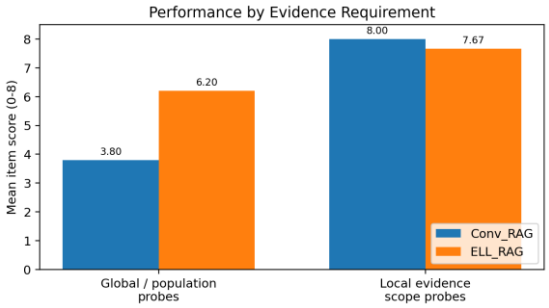


Figure 3: Mean scores on the prespecified global/population-level and local-evidence task groupings.

Figure 3 summarizes the central design pattern: the architectures were nearly equivalent when evidence was local and directly retrievable, whereas ELL_RAG showed a substantially larger descriptive advantage when the task required global/population-level measurement evidence. The near-tie on local questions is important because it indicates that the measurement-oriented extension did not create a uniform advantage; conventional retrieval remained highly capable when its evidence mechanism matched the task.

3.3 RQ3. Architectural Strengths and Limitations

ELL_RAG: ELL_RAG's principal strength was connecting population-level questions to full-dataset measurement evidence, including reliability statistics, descriptive summaries, and controlled comparisons. Its main limitation was orchestration: appropriate tools could be available yet incompletely selected, executed, or combined. Q07, Q13, and Q14 illustrate that measurement-oriented RAG requires reliable routing, evidence verification, and multi-source composition in addition to deterministic functions.

Conv_RAG: Conv_RAG's principal strength was local, directly retrievable evidence. It performed at or near ceiling on several rubric-,

metadata-, and professional-interpretation items and frequently abstained appropriately when dataset-level evidence was unavailable. Its main limitation was structural: without a deterministic analysis pathway, it often could not answer full-dataset reliability or subgroup questions. For example, Q13 showed consequential risk of using locally retrieved examples in place of a requested population-level analysis.

3.4 Synthesis Across the Three Research Questions

Across the three research questions, the findings point to task–evidence alignment rather than a simple comparative ranking of the two architectures. ELL_RAG's advantage was concentrated in tasks requiring deterministic computation or dataset-level aggregation, whereas conventional retrieval remained highly effective when the relevant evidence was local and directly accessible. The results therefore support a selective extension of conventional RAG with measurement-oriented capabilities, rather than broad replacement of the conventional architecture, while also highlighting the additional routing and evidence-integration demands introduced by tool-augmented systems.

4 Discussion

The findings support a design-oriented interpretation in which the value of measurement-oriented RAG depends on evidentiary demand. ELL_RAG showed its clearest descriptive advantage on questions requiring global or population-level evidence, whereas the two systems performed nearly equivalently when the needed evidence was local and directly retrievable. This pattern is more informative than the overall score difference because it shows that ELL_RAG's advantage was concentrated in tasks requiring dataset-level computation rather than retrieval alone. It also distinguishes retrieval relevance from evidentiary sufficiency: a retrieved essay may be highly relevant to a question but still be insufficient to support subgroup or population-level conclusions. This distinction helps explain the complementary strengths of the two architectures. Conventional RAG remained effective when relevant evidence could be retrieved directly, while measurement-oriented RAG provided additional support when the requested inference required structured computation, aggregation, or analysis across the broader dataset.

At the same time, the findings show that additional analytical capabilities do not automatically produce stronger responses. Their value depends on appropriate routing, accurate execution, effective evidence integration, and verification. In this sense, tool orchestration is a central design requirement rather than a secondary implementation detail. Appropriate acknowledgment of insufficient evidence is equally important because a system should state when the available evidence cannot support the requested statistic or conclusion, rather than provide an unsupported quantitative answer.

The observed pattern is consistent with broader RAG research distinguishing localized retrieval from tasks requiring corpus-level aggregation or structured-data access (Edge et al., 2024; Murtaza et al., 2025). The present study does not establish a universal local-versus-global rule, but it extends this distinction into measurement-oriented assessment analytics. The practical implication is therefore not to replace conventional RAG broadly, but to match the evidence mechanism to the level of inference required: use retrieval when retrieval is sufficient, invoke deterministic analysis when dataset-level evidence is needed, coordinate multiple evidence sources when necessary, and retain human oversight for consequential interpretations.

5 Limitations

These findings remain preliminary because the study used a small prespecified question set and one author-rater. The results should therefore be interpreted as descriptive evidence about architectural behavior rather than as proof of general system superiority. The sampled questions were designed to probe measurement-oriented behavior and do not represent the full range of possible RAG tasks. Future work should examine whether the same local-versus-global pattern replicates with broader task sampling and independent raters. Further development should also strengthen routing, multi-source evidence integration, and post-tool verification.

6 Conclusion

The study supports a bounded design principle: retrieval, computation, and interpretation should be aligned with the evidentiary level of the question**.** Conventional RAG remains effective when evidence is local and directly retrievable, whereas measurement-oriented extensions appear most useful when the task requires structured computation, aggregation, or population-level analysis. The contribution is therefore not a claim of universal architectural superiority, but a selective design approach in which additional measurement-oriented capabilities are introduced when the evidentiary demands of the task justify them.

References

American Educational Research Association, American Psychological Association, & National Council on Measurement in Education. (2014). *Standards for educational and psychological testing*. American Educational Research Association.

Cohen, J. (1968). Weighted kappa: Nominal scale agreement with provision for scaled disagreement or partial credit. *Psychological Bulletin, 70(4)*, 213-220.

Crossley, S., Tian, Y., Baffour, P., Franklin, A., Kim, Y., Morris, W., Benner, M., Picou, A., & Boser, U. (2023). The English Language Learner Insight, Proficiency and Skills Evaluation (ELLIPSE) Corpus. *International Journal of Learner Corpus Research, 9*(2), 248–269.

Edge, D., Trinh, H., Cheng, N., Bradley, J., Chao, A., Mody, A., Truitt, S., Metropolitansky, D., Ness, R. O., & Larson, J. (2024). From local to global: A graph RAG approach to query-focused summarization [Preprint]. arXiv. https://doi.org/10.48550/arXiv.2404.16130

Gao, Y., Xiong, Y., Gao, X., Jia, K., Pan, J., Bi, Y., Dai, Y., Sun, J., Wang, M., & Wang, H. (2023). *Retrieval-augmented generation for large language models: A survey*. arXiv. https://arxiv.org/abs/2312.10997

Kane, M. T. (2013). Validating the interpretations and uses of test scores. *Journal of Educational Measurement, 50*(1), 1–73.

Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W., Rocktäschel, T., Riedel, S., & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. arXiv. https://arxiv.org/abs/2005.11401v4.

Messick, S. (1989). Validity. In R. L. Linn (Ed.), *Educational Measurement* (3rd ed., pp. 13–103). American Council on Education / Macmillan.

Murtaza, S. S., Nie, Y., Avan, E., Soni, U., Liao, W., Carnegie, A., Mathias, C. J., Jiang, J., & Wen, E. (2025). Implementing retrieval augmented generation technique on unstructured and structured data sources in a call center of a large financial institution. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 3: Industry Track)* (pp. 598-606). Association for Computational Linguistics. https://doi.org/10.18653/v1/2025.naacl-industry.48

Pack, A., Barrett, A., & Escalante, J. (2024). Large language models and automated essay scoring of English language learner writing: Insights into validity and reliability. *Computers and Education: Artificial Intelligence, 6*, 100234.

A Design Workflow for Measurement-Oriented ELL_RAG

ELL Writing Dataset

- 5,840 complete cases
- Essay text, scores, rubrics, and rater information
- Grade, Ethnicity, Gender, Economically Disadvantaged Status, and Prompts
- Linguistic features: num_words, MTLD, and TTR

Deterministic Measurement-Tool Construction

- Group and full-dataset descriptive summaries
- Controlled subgroup comparisons
- Reliability metrics
- Rater-disagreement analyses
- Rubric lookup
- Dataset-level linguistic-feature summaries
- Identification of large rater-disagreement cases

Tool-Level Functional Checks

- Check intended analytic-tool invocation
- Review structured statistical outputs

- Check application of grade and prompt controls where applicable
- Check warnings for multiple prompts and small samples
- Check unsupported or unavailable comparisons
- Functional verification of tools before formal system comparison

Essay Vectorization and Retrieval Preparation

- Convert 5,840 essays into LangChain Documents
- Store score and metadata with each essay
- Embed essays using sentence-transformers/all-MiniLM-L6-v2
- Store embedded essays in Chroma vector database
- Confirm collection count of 5,840 documents

Preliminary Semantic-Retrieval Check

- Execute targeted similarity queries
- Inspect retrieved essay text and associated metadata
- Review similarity results for consistency with intended query characteristics
- Functional retrieval check only; not a formal retrieval-validation study

User Measurement Question

Evidence-Demand Classification and Routing

- Local Evidence Route
  - Retrieve relevant essays when local examples are requested
  - Use official rubric lookup for score-meaning questions
  - Use structured metadata for eligibility or factual constraints
  - Treat retrieved essays as local illustrative or diagnostic evidence only
- Global / Population-Level Measurement Route
  - Use full-dataset descriptive summaries when corpus-level quantities are requested
  - Use reliability and rater-disagreement functions for rater-level questions
  - Use controlled subgroup functions for group-level comparisons
  - Use dataset-level linguistic summaries for corpus- or subgroup-level feature questions
  - Do not infer population quantities from a small set of retrieved essays
- Combined Evidence Route
  - Combine deterministic measurement evidence with rubric and/or local essay evidence when needed
  - Preserve the distinct evidentiary role and scope of each source
  - Add sample-size, scope, and interpretation cautions where appropriate

Evidence Assembly

- Dataset-level measurement evidence
- Official rubric evidence
- Reliability and controlled-comparison evidence
- Structured metadata
- Local essay evidence
- Dataset-level linguistic-feature evidence

LLM Synthesis Layer

- GPT-4o with temperature = 0.0
- LangChain-based orchestration
- Uses retrieved evidence and deterministic tool outputs
- Does not independently calculate dataset-level statistics
- Does not define official rubric meanings without rubric lookup
- Maintains local-versus-global evidence boundaries during synthesis

Measurement-Oriented Response

- Direct interpretation or summary
- Relevant measurement evidence
- Evidence scope or contextual classification, when needed
- Caution and limitations
- Sources / evidence identifiers

B Representative Evaluation Questions

Category 1: Validity, fairness, and quality in assessment interpretation

Question: Suppose the official Cohesion rubric descriptor for score 3 is retrieved correctly, full-dataset Cohesion rater reliability is available, and several retrieved essays with Cohesion score 3

appear broadly compatible with the descriptor. What interpretation is supported when these sources are considered together, and what broader validity or fairness claim would still require additional evidence?

Category 2: Subgroup fairness, transparency, and human oversight

Question: For Grammar, examine rater agreement separately by observed economic status category. Report paired sample size, exact agreement, agreement within one point, and weighted kappa where estimable, and identify any subgroup result that should be treated cautiously because of limited data or uncertainty.

Category 3: Integration of measurement and retrieval evidence

Question: A retrieved essay assigned Vocabulary score 3 appears weaker than the official Vocabulary score-3 descriptor might suggest. State the score meaning from the official rubric and explain how the apparently inconsistent essay should—and should not—be used as local evidence.

Category 4: Responsible score interpretation and professional use

Question: A teacher asks the system to replace a student's official Conventions score because several semantically similar essays received higher scores. How should the system respond, and what student-specific evidence, rubric-based review, and authorized human process would be required for a legitimate score challenge?

Category 5: Ethical handling of evidence, uncertainty, and overclaiming

Question: Compare economic-status group patterns in num_words, TTR, and MTLD using dataset-level linguistic summaries. Describe the largest observed differences, what they may suggest about the sampled writing, and which broader conclusions are and are not supported by these features alone.

Note. One representative question is shown from each of the five prespecified evaluation categories. Questions are provided to illustrate the scope of the evaluation and were not selected based on system performance.

C Illustrative Paired Responses from Conv_RAG and ELL_RAG

Question: For Grammar, examine rater agreement separately by observed economic status category. Report paired sample size, exact agreement, agreement within one point, and weighted kappa where estimable, and identify any subgroup result that should be treated cautiously because of limited data or uncertainty.

Conv_RAG Response

Summary

The analysis of rater agreement for the Grammar score by observed economic status category involves examining paired sample size, exact agreement, agreement within one point, and weighted kappa. However, the retrieved context does not provide specific numerical results or detailed subgroup analyses for these metrics.

Evidence

The corpus documentation indicates that the dataset includes various writing-score fields, including Grammar, and contextual fields such as economic status (SES) [Source 8]. However, specific data on rater agreement metrics like exact agreement, agreement within one point, and weighted kappa by economic status category is not provided in the retrieved context.

Interpretation

In general, rater agreement metrics such as exact agreement, agreement within one point, and weighted kappa are used to assess the consistency and reliability of ratings across different raters. Exact agreement measures the percentage of cases where raters give the same score. Agreement within one point considers scores that are within one point of each other as agreeing. Weighted kappa accounts for the degree of disagreement, providing a more nuanced measure of agreement that considers the severity of disagreements.

For subgroup analyses, such as by economic status, it is important to ensure that there is sufficient data for each subgroup to produce reliable estimates. Small sample sizes can lead to high uncertainty in agreement metrics, making some subgroup results less reliable.

Limitations

The retrieved context does not provide specific numerical results or detailed analyses of rater agreement metrics by economic status category. Without this information, it is not possible to report the paired sample size, exact agreement, agreement within one point, or weighted kappa for each economic status category. Additionally, without subgroup-specific data, it is not possible to identify which subgroup results should be treated cautiously due to limited data or uncertainty.

Sources

[Source 8]: ELLIPSE analytic dataset documentation.

ELL_RAG Response

Reliability Interpretation

The reliability of Grammar scoring was examined separately for economically disadvantaged and not economically disadvantaged groups. The analysis focused on exact agreement, agreement within one point, and weighted kappa for each subgroup.

Rater Agreement Evidence

For the economically disadvantaged group (n = 4050), the exact agreement between raters was 51.3%, with 99.0% agreement within one point. The weighted kappa was 0.556, indicating moderate agreement. For the not economically disadvantaged group (n = 1789), the exact agreement was slightly higher at 52.1%, with 99.4% agreement within one point. The weighted kappa for this group was 0.628, suggesting a moderate-to-strong agreement.

Essay Context, Not Reliability Evidence

No essay-level diagnostic evidence was used in this analysis. Therefore, specific essay features that might have influenced rater agreement are not available for contextual interpretation.

Caution About Sample Size and Generalization

Both subgroups met the minimum paired sample size required for reporting reliability coefficients. However, while the sample sizes are sufficient for reporting, the reliability indices should be interpreted with caution as they do not establish precision or stability of the subgroup estimates. Additionally, subgroup agreement differences are descriptive signals for further review and do not prove scoring bias or fairness.

Sources: No specific text_ids were used in this analysis.

Note. Response headings may vary according to the evidentiary demands of the question. Responses are reproduced verbatim from the frozen evaluation run.

Automated item evaluation: Predicting item acceptance and rejection using LLM-generated critiques

Hotaka Maeda [1,2]*
Yikai Lu [3]

[1]Smarter Balanced, Santa Cruz, CA, USA
[2]University of California, Santa Cruz, Santa Cruz, CA, USA
[3]Department of Educational Psychology, University of Minnesota-Twin Cities, Minneapolis, MN, USA

Abstract

We built a near-comprehensive automated item evaluation model, predicting historical item acceptance or rejection from item text and Qwen3-generated critiques using 52,759 items from a large-scale testing program. Two DeBERTaV3 classifiers fused reached AUC .80 overall and .86 for math. Fairness-related rejections remained difficult, underscoring the need for human review.

1 Introduction

Evaluating the quality of new items is expensive and time-consuming, requiring content review and field testing. This issue has grown since the emergence of automated item generation (AIG) using large language models (LLMs), and AIG review papers consistently urge for methods that can quickly filter poor quality items (Circi et al., 2023; Falcão et al., 2022; Tan et al., 2025). Relevant literature has focused on using LLMs, natural language processing (NLP), or machine learning to predict individual qualities of items, such as item difficulty (AlKhuzaey et al., 2024), content alignment with the assessment blueprint (Fu et al., 2025), or differential item functioning (Maeda and Lu, 2025). To define this broader problem space, we use the term *automated item evaluation* (AIE; Yu and Burke, 2026), extending its prior use from a narrower scope. Here, AIE refers to the use of computational methods to automatically assess item quality without requiring manual expert review or field testing of the items under evaluation.

The difficulty of AIE stems from the countless reasons an item can be unsuited for operational use. The Standards for Educational and Psychological Testing (AERA et al., 2014) articulate the requirements that define a valid, reliable, and fair item, and Haladyna and Downing (1989) cataloged 43 item-writing rules covering stem construction, correct answer selection, and distractor quality. Training and deploying a separate prediction model for each rule is unrealistic, as model training requires substantial sample sizes and computational resources. Detecting item flaws separately may also be inefficient for rejecting undesirable items.

A comprehensive AIE approach is therefore desirable. One possibility is to use an existing item bank, treating operational items as accepted" and items permanently retired due to irreparable issues as rejected", then training a model to predict that outcome from item content alone. Accepted items had to pass rigorous content expert review and field testing, so this single label becomes an all-encompassing measure of item quality. Although this approach requires a large established item bank, it may be the single most comprehensive AIE method for both AIG and human-written items.

To our knowledge, no studies have predicted item acceptance as an AIE method. We train a transformer language model to classify items as accept or reject for operational use, using item status data from a large-scale standardized testing program, while augmenting the data using item critiques generated by an LLM. Model behavior is analyzed using historical item developer comments and sentiment analysis.

2 Related Works

2.1 Item Evaluation

Item evaluation is an important step during test development (AERA et al., 2014; Haladyna and Downing, 1989). Gorgun and Bulut (2025) summarized existing approaches into three categories: metric-based evaluations, post-hoc analysis, and human evaluations. Metric-based approaches typically compare items with a set of reference items,

*Correspondence concerning this article should be addressed to Hotaka Maeda, Smarter Balanced, University of California-Santa Cruz, 1156 High St, Santa Cruz, CA 95064. E-mail: hotaka.maeda@smarterbalanced.org

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 113–129
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

including NLP-based techniques such as word count, TF-IDF (Benedetto et al., 2020), cosine similarity for identifying enemy items, and measures of readability and text complexity (Yu and Burke, 2026; Amini et al., 2025). Pelánek et al. (2022) used an outlier detection mechanism to flag statistically deviant items. However, metric-based approaches are limited by their reliance on reference items, which can restrict the content domain of generated items, as valid items may be penalized simply because they *appear* different from reference items at the level of linguistic features (Gorgun and Bulut, 2025).

Post-hoc analysis approaches instead involve field testing to collect item response data from a representative sample of examinees. The Standards recommend meeting psychometric specifications related to item difficulty, item discrimination, and the absence of differential item functioning (AERA et al., 2014). However, this is costly and time-consuming, requiring a substantial sample size for reliable estimates. Finally, human evaluators can identify problems difficult to detect computationally, but human evaluation is subjective and time-consuming, and finding suitable experts is challenging.

In summary, metric-based approaches are automatic but limited in scope, while post-hoc analyses and human evaluations can be high-quality but resource-intensive. We argue in the next section that AIE can combine the advantages of all three while bypassing some of their limitations.

2.2 "Automated" Item Evaluation

The term "automated item evaluation" or similar has been used occasionally, typically in the AIG literature (Gorgun and Bulut, 2025; Prentzas and Binopoulou, 2025; Yu and Burke, 2026; Shin and Gierl, 2024; Wang et al., 2025). Despite sparse usage and variation in scope, definitions of AIE share a common theme: the use of computational methods to automatically assess some facets of item quality without requiring manual expert review or field testing. Among the three categories described by Gorgun and Bulut (2025), only the metric-based approach meets this criterion, as it eliminates reliance on field testing or human evaluators.

That said, results from post-hoc analyses or human evaluations could serve as labels for training AIE prediction models. Gorgun and Bulut (2025) trained an LLM to identify low-quality AIG items, but such an approach provides only a coarse-grained evaluation, classifying items simply as *good* or *bad*. Pelánek et al. (2022) argues this may be a fundamentally ill-defined problem, as items warranting attention depend on context and are partly subjective. Items may also be problematic for many reasons, including linguistic or psychometric characteristics that human evaluators cannot always capture, and vice versa. Predicting statistical or psychometric properties from past post-hoc data has received more attention, including item difficulty (AlKhuzaey et al., 2024; Benedetto et al., 2023), item discrimination (e.g., Maeda and Lu, 2026; Han et al., 2025; Yaneva et al., 2020), and differential item functioning (Maeda and Lu, 2025). Although linguistic information was used as input, the outcomes were limited to psychometric characteristics rather than directly assessing linguistic or semantic dimensions of item quality.

LLM evaluators could be considered a fourth category. LLMs can mimic human evaluators but are automated, and unlike fine-tuning they require no labeled training data and can run zero-shot. Notably, LLM-based evaluation has enabled automated assessment of item-content alignment with assessment blueprints, an area that has received limited attention (Fu et al., 2025), suggesting LLM-based evaluation could shed light on aspects difficult to evaluate automatically.

The major limitation of the metric-based and psychometric-prediction approaches is that they often focus on one aspect of item quality at a time, even though different dimensions may be interrelated. For example, a psychometrically difficult mathematics item may also contain difficult words, a relationship that cannot be leveraged when these issues are evaluated separately. Relying solely on NLP or psychometric evaluations may also overlook problems identifiable only through human judgment. It is therefore preferable to construct a model that evaluates multiple dimensions simultaneously. Our comprehensive approach does so by using a dataset of rejected items with a broad range of documented rejection reasons. Rather than relying solely on human evaluators, psychometric analyses, or linguistic features, we fine-tune a Transformer language model to distinguish acceptable from unacceptable items while implicitly learning from multiple sources of item-quality evidence.

2.3 Transformer Language Models for Text Classification

The rise of Transformer language models increased demand for AIE research by making large-scale AIG feasible, while expanding the computational tools useful for AIE. Traditionally, text-based prediction in assessment relied on expert judgment (Wauters et al., 2012), syntactic features such as word count and term frequency (Benedetto et al., 2020), or semantic features like word embeddings (Hsu et al., 2018). More recently, Transformer-based models have substantially improved accuracy (Li et al., 2025; Han et al., 2025; Maeda, 2025). Introduced by Vaswani et al. (2017), the Transformer replaced recurrent networks with parallel attention, enabling efficient training and scalable modeling of long-range dependencies. Models such as BERT (Devlin et al., 2019) are pre-trained on large corpora and fine-tuned for downstream tasks, with text tokenized into embeddings and passed through encoder layers to produce contextually enriched representations.

We use DeBERTaV3-large (He et al., 2021), which extends BERT (Devlin et al., 2019) and RoBERTa (Liu et al., 2019). DeBERTa disentangles content and positional information within its attention mechanism, representing each token by separate content and position vectors to capture relationships between words more effectively. The V3 variant adds replaced token detection, yielding a 304-million-parameter model whose natural language understanding even outperforms the newer ModernBERT (Warner et al., 2024).

2.4 Using LLM Critiques for AIE

Recently, language models have been used to extract reasoning or rationales to augment a task-specific model (Henrichsen and Krebs, 2025; Hsieh et al., 2023; Scarlatos et al., 2025). In the assessment context, GPT-4o (OpenAI, 2024) was used to generate the reasoning steps required to reach each multiple choice option, which were then input into the Longformer encoder to predict item difficulty (Feng et al., 2025). In another example, LLMs were given the role of content experts to make absolute and pairwise item comparisons to estimate difficulty (Kolesnikova et al., 2026).

We use a similar data augmentation technique with Qwen3 (Qwen Team, 2025). The Qwen3 follow the broader trend of capable small language models (SLMs) suited to resource-constrained settings (Lu et al., 2025). A key capability distinguishing them from earlier SLMs is the integration of a *thinking mode* and a *non-thinking mode* within a single model, letting users toggle between slow, multi-step reasoning and fast response generation at inference time. Empirical results in Qwen Team (2025) indicate that the Qwen3-8B, 4B, and 1.7B base models each outperform Qwen2.5 models of the next larger size class.

3 Methods

3.1 Item Data

The study included 52,759 items designed for English language arts (ELA) and mathematics state summative assessments for grades 3–11. Of these, 34% had been rejected, defined as permanently removed from future operational use (30% for ELA, 39% for math). Rejection can occur at any stage of item development, before or after field testing, or after some operational use, and spans a range of reasons including content and standards alignment, psychometrics, accessibility and sensitivity, and non-content issues such as data integrity problems (see Appendix A). Items that had not been rejected and had not yet been field tested were excluded, as their operational suitability remains unknown. Items rejected prior to field testing were retained, as their unsuitability is established (14% of rejected items). Items were field-tested among grade-matched students across multiple states in the United States.

To prepare for modeling, item text was concatenated with a `\n` separator between prompts and any available answer options. Selection-based (e.g., multiple choice, multiple select) and constructed response (e.g., short answer, extended response) items were included. Reading or listening passages were excluded due to their extensive length. Images, figures, tables, metadata, scoring rubrics, answer keys were also excluded. Data were partitioned randomly into approximately 80% training, 10% validation, and 10% test, with multiple items associated with the same stimulus always kept in the same group.

3.2 Prediction Approaches

Five prediction approaches were compared. All model training and inference used a single NVIDIA A10G Tensor Core 24GB graphics processor.

3.2.1 Zero-Shot Baseline

Qwen3-0.6B (Qwen Team, 2025) was prompted to classify each item as "accept" or "reject" directly, without task-specific fine-tuning (see Appendix B for the prompt). The subject, grade level, and aligned common core standards were included in the prompt.

3.2.2 Raw Item Text Fine-Tuned Model

Item text was tokenized and passed into DeBERTaV3-large (He et al., 2021), fine-tuned as a binary classifier using Low-Rank Adaptation (LoRA; Hu et al. 2022; rank = 32, alpha = 64, LoRA dropout = 0.05, dropout = 0.1, batch size = 8, max token length = 512, weight decay = .01, learning rate = 5e-5, epochs = 3). LoRA is a parameter-efficient method that inserts trainable low-rank matrices into selected layers while keeping pretrained weights frozen, substantially reducing memory and computational costs. We added a binary classification head with a sigmoid function to convert logits into rejection probabilities, followed by cross-entropy loss (CEL).

Accepted items outnumbered rejected items roughly two to one. This imbalance was addressed with class weights in the CEL, computed inversely proportional to class frequency so the minority class received greater influence on gradient updates. This preserves the full training sample while discouraging the model from defaulting to the majority class. We retained the state with the lowest validation CEL across all epochs.

3.2.3 Item Critique Fine-Tuned Model

To augment the raw item text, Qwen3-0.6B (Qwen Team, 2025) with thinking mode enabled was prompted to generate a two-sentence critique or praise of each item's quality, with subject, grade level, and aligned common core standards included in the prompt (see Appendix C for the prompt and Appendix D for example critiques). The resulting critique text was used as input to DeBERTaV3-large in place of the raw item text, with the same procedure and hyperparameters as the text-only model except LoRA rank = 50 and alpha = 100.

3.2.4 Text + Critique Fusion Model

The fusion model combines representations from both DeBERTaV3-large models above. The 1024-dimensional final-layer output of each was extracted and concatenated. The DeBERTa weights were frozen and only the fusion layers trained: a 128-dimensional intermediate layer followed by a binary classification head with a sigmoid function and a weighted CEL (see Figure 1).

3.2.5 Stand-Alone Subject Models

Because ELA and math are distinct subjects, we also re-trained separate raw text-only, critique-only, and fusion models for each.

3.3 Evaluation

The area under the ROC curve (AUC) measured overall discriminative ability independent of the cutoff threshold. With the classification threshold fixed at .50 by default, we report accuracy (proportion correctly classified), precision (proportion of flagged items truly rejected), sensitivity (proportion of rejected items correctly flagged), specificity (proportion of viable items correctly passed), and F1 (harmonic mean of precision and sensitivity). Rejected cases are treated as positive.

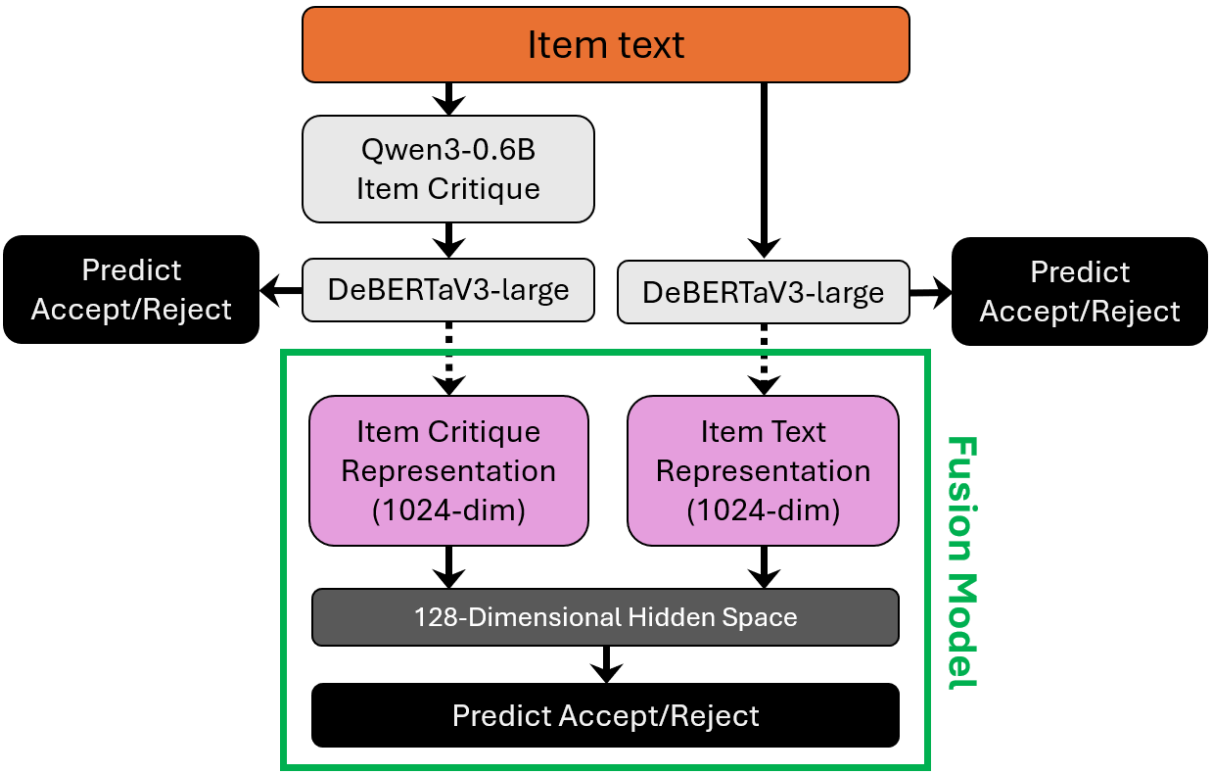
3.3.1 Rejection Reason Classification

Item developers often left comments when moving an item to permanent rejection, though quality varied from absent, uninformative (e.g., "Moving to rejected"), or detailed. Multiple comments for the same item were concatenated into a single record. Using Claude Sonnet 4.6 (Anthropic, 2025) with manual fixes, comments from rejected items were cleaned and classified into 10 rejection reason categories: content, psychometric, data review, bias, abandoned, incomplete, passage, scoring, non-content, and no data (see Appendix A for definitions). These labels, used only for test sets, assessed how well the best-performing model detected items across each rejection reason. Only 19% of rejected test items had usable rejection reason data. The rest fell into "no data" because comments were missing (80%) or ambiguous (1%).

3.4 Rejection Rate by Difficulty and Discrimination

We examined Spearman correlations between fusion model predictions and several difficulty and discrimination metrics, including proportion correct, item-total correlation, and item response theory (IRT) parameters. All items were calibrated using the two-parameter logistic (2PL) model (Birnbaum, 1968) for dichotomous items or the generalized partial credit model (GPCM; Muraki 1992) for polytomous items. For GPCM items, the difficulty parameter was centered to represent the average

Figure 1: Automated item evaluation model using raw text + LLM Critique fusion.



Note. Item text is processed directly through the first Deberta. Qwen is used to generate item critiques, then entered into a separate Deberta. Two DeBERTas are combined in the final fusion model.

difficulty of the thresholds, so both discrimination (a) and difficulty (b) were comparable between 2PL and GPCM. Although the data spanned two subjects and multiple grades, we standardized θ to mean zero and standard deviation one within each subject-grade combination, removing the vertical scale for this analysis. All statistics, including proportion correct and item-total correlation, are therefore relative to subject and grade. For example, a low proportion correct for an ELA grade 3 item means it is difficult for grade 3 students. The 6.4% of items rejected prior to field testing were excluded.

3.5 Sentiment Analysis of Critiques

To examine whether sentiment partly explains how the Qwen3 critiques improved prediction, we computed sentiment using SiEBERT (Hartmann et al., 2023). This is a RoBERTa-large model fine-tuned for binary sentiment classification across 15 heterogeneous English text sources. For each critique, we extracted the difference between the positive and negative class logits, where higher values indicate more positive sentiment. We computed Spearman correlations between critique sentiment and the predicted rejection probability of the text-only model P_{text}, the fusion model P_{fusion}, and their difference $P\Delta = P_{fusion} - P_{text}$.

4 Results

All results are based on test data. The zero-shot approach had the lowest F1 score (F1 = .23; see Table 1). The fusion model consistently outperformed both the raw text-only and critique-only models (F1 = .64). Prediction for math was considerably more accurate than ELA in all cases, and the standalone subject models did not meaningfully change performance. A fusion model with non-weighted CEL performed essentially the same (F1 = .63), but the weighted model offers an interpretive advantage: class weighting shifts the optimal decision boundary near 0.5, aligning the classification threshold with the intuitive probability midpoint. Subsequent analyses evaluate the fusion model.

4.1 Improving Sensitivity by Lowering the Cutoff Threshold

Lowering the cutoff below .5 increases sensitivity to flag non-viable items at the cost of specificity (see Figure 2). A threshold of .25 raises overall sen-

Table 1: Classification Performance by Model and Subject

Model	p	AUC	Accuracy	Precision	Sensitivity	Specificity	F1
Zero-shot	.17	.50	.61	.35	.17	.83	.23
Raw text-only	.31	.76	.73	.61	.56	.82	.59
ELA	.21	.69	.72	.55	.39	.86	.46
Math	.42	.80	.74	.65	.71	.76	.68
Critique-only	.30	.72	.71	.59	.52	.82	.55
ELA	.20	.63	.69	.49	.32	.85	.39
Math	.41	.79	.74	.65	.68	.77	.67
Fusion (raw text + critique)	.34	.80	.75	.63	.64	.81	.64
ELA	.27	.72	.73	.54	.49	.83	.51
Math	.43	.86	.78	.70	.77	.79	.73
ELA Fusion (stand-alone)	.33	.69	.69	.47	.53	.75	.50
Math Fusion (stand-alone)	.39	.85	.79	.72	.73	.82	.73

Note. ELA = English language arts, p = proportion of items predicted as reject, AUC = area under the curve. Cutoff threshold for models were fixed to .5.

sitivity to .90 while reducing specificity to .42 and F1 to .60. Math reaches sensitivity .91 (specificity = .56, F1 = .70) and ELA .88 (specificity = .31, F1 = .51; see Appendix E).

4.2 Item Rejection Reasons

Predicted rejection probability by reason was estimated on truly rejected test items (see Figure 3). Abandoned items (sensitivity $=$ .75) and those with psychometric issues (sensitivity $=$ .66) had among the highest sensitivities, especially for math. Compared to the text-only model, incorporating critiques improved detection across nearly all categories, substantially so for incomplete math items (Δ sensitivity $=$ .31) and content issues (Δ sensitivity $=$ .33). Low sensitivity (.16) for non-content rejections was expected, as they typically reflect back-end data integrity issues unrelated to item text. However, the fusion model struggled with bias, sensitivity, fairness, and accessibility concerns (sensitivity $=$ .27), 91% of which were ELA.

4.3 Item Difficulty and Discrimination

The model assigned higher rejection probabilities to items with lower proportion correct ($r = -.30$). Correlations with total item correlation ($r = -.06$), IRT discrimination ($r = -.07$), and the intercept ($-1.7ab$; $r = -.19$) were weaker. Overall, the model is sensitive to rejecting items that may be difficult for the target population (see Figure 4).

4.4 Critique Sentiment

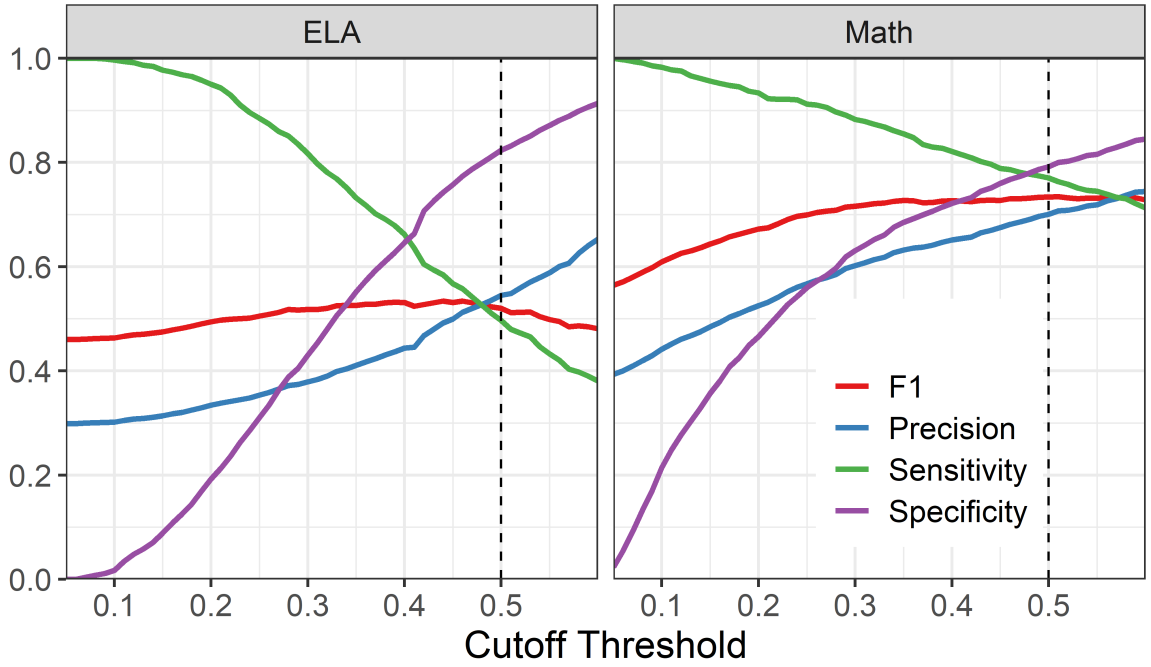
Correlations between critique sentiment and P_{text} and P_{fusion} were significant only for math P_{fusion} ($r = -.05$, $p < .05$). Negative critique sentiment was associated with slightly increased $P\Delta$ across all subject and status groups. Correlation between sentiment and $P\Delta$ was $-.05$ for ELA ($p < .01$) and $-.14$ for math ($p < .0001$). Disaggregating by status, correlations were not significant for ELA ($p > .05$), and were $-.11$ for accepted and $-.18$ for rejected math items ($p < .0001$). The association was thus stronger for math than ELA, and strongest among truly rejected math items. Negative sentiment increasing rejection rates was expected, but sentiment was overall a minor characteristic of how critiques related to model behavior.

5 Discussion

We set out to build a comprehensive AIE model that could identify unwanted items across as many aspects of item quality as possible. We achieved this by training a transformer language model to classify items as accepted or rejected for operational use, using item status data from a large-scale standardized testing program. Data augmentation using item critiques generated by Qwen3 improved prediction accuracy across nearly all rejection reasons. The result was a single, near-comprehensive AIE model, more convenient than most other AIE methods that target a single aspect of item quality, with practical promise for large-scale assessment pipelines where both the evaluation and discard of items carry substantial cost.

AIE is useful for both human-written and AI-generated items, though the purpose and strategy of its use can vary. In AIG contexts, where item generation is inexpensive, a lower detection threshold such as .25 may be preferable. At the cost of specificity (.42) and F1 (.60), sensitivity of our

Figure 2: Cutoff threshold analysis for fusion model (raw text + critique)



Note. ELA = English language arts. Lowering the cutoff threshold below .5 increases sensitivity to flag rejected items at the cost of specificity. A lower cutoff threshold may be preferable in contexts such as automated item generation where item generation has a low cost.

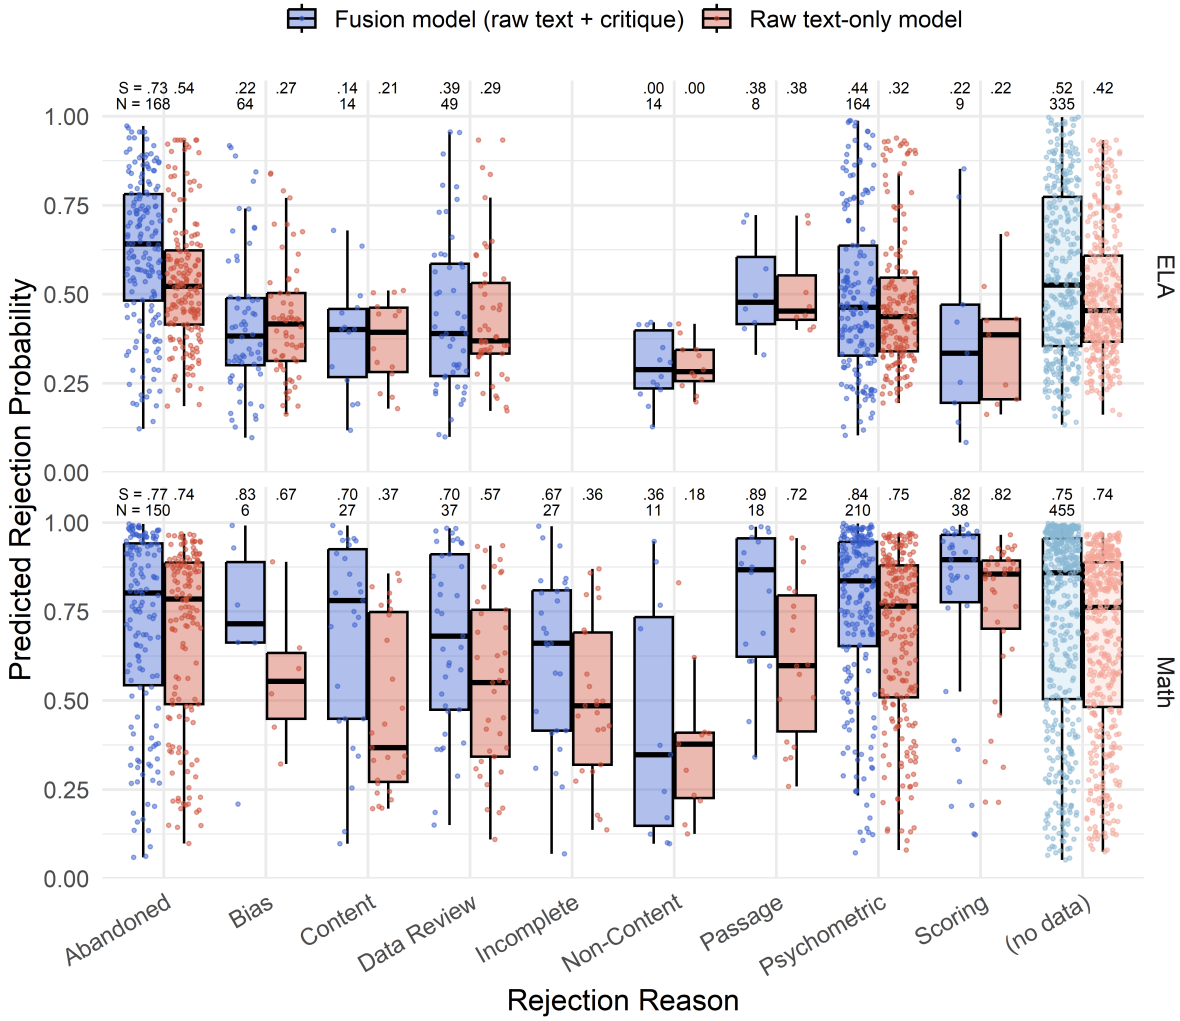
fusion model rises to .90, catching most future rejections before they happen and saving substantial content expert review time. In more traditional item development contexts, the default threshold of .5 may serve as an effective early warning to revise the item. Predictions were considerably more accurate for math than ELA, which may reflect fundamental differences in how items are rejected between subjects. In particular, ELA rejections were often related to their passages, which were excluded from the current study.

Associations with item statistics showed that the model rejected items that were too difficult. This shows the value of item difficulty prediction research (AlKhuzaey et al., 2024; Benedetto et al., 2023), but the modest correlation (r=-.30) suggests difficulty prediction alone is insufficient. Rejection rates were nearly unrelated to item discrimination, showing that items were rejected for many reasons beyond difficulty and discrimination.

Several aspects of this study are novel. First, fusing an encoder representation of raw item text with an encoder representation of a separate decoder-generated critique is relatively new, though a handful of studies have taken similar approaches (Henrichsen and Krebs, 2025; Hsieh et al., 2023; Scarlatos et al., 2025; Feng et al., 2025). This method leverages the decoder's capacity to draw on broad world knowledge when generating critiques, while relying on the encoder for the precise, task-specific prediction that decoders are not designed to produce. The approach could likely improve with a larger decoder, as our Qwen3-0.6B model is very small (600 million parameters) compared to proprietary multimodal models rumored to have trillions of parameters (Li, 2026), though such models may require additional security measures unlike locally run models like Qwen. Second, we are not aware of any other paper that has used item developer comments to help build a prediction model in educational assessment. The closest precedent may be (Ma, 2025), who applied DistilBERT and machine learning to examinee comments to identify those most relevant for item review, establishing that free-text commentary contains information for item quality decisions. We considered using item developer comments in training but decided against it, as only 19% of our data had rejection comments.

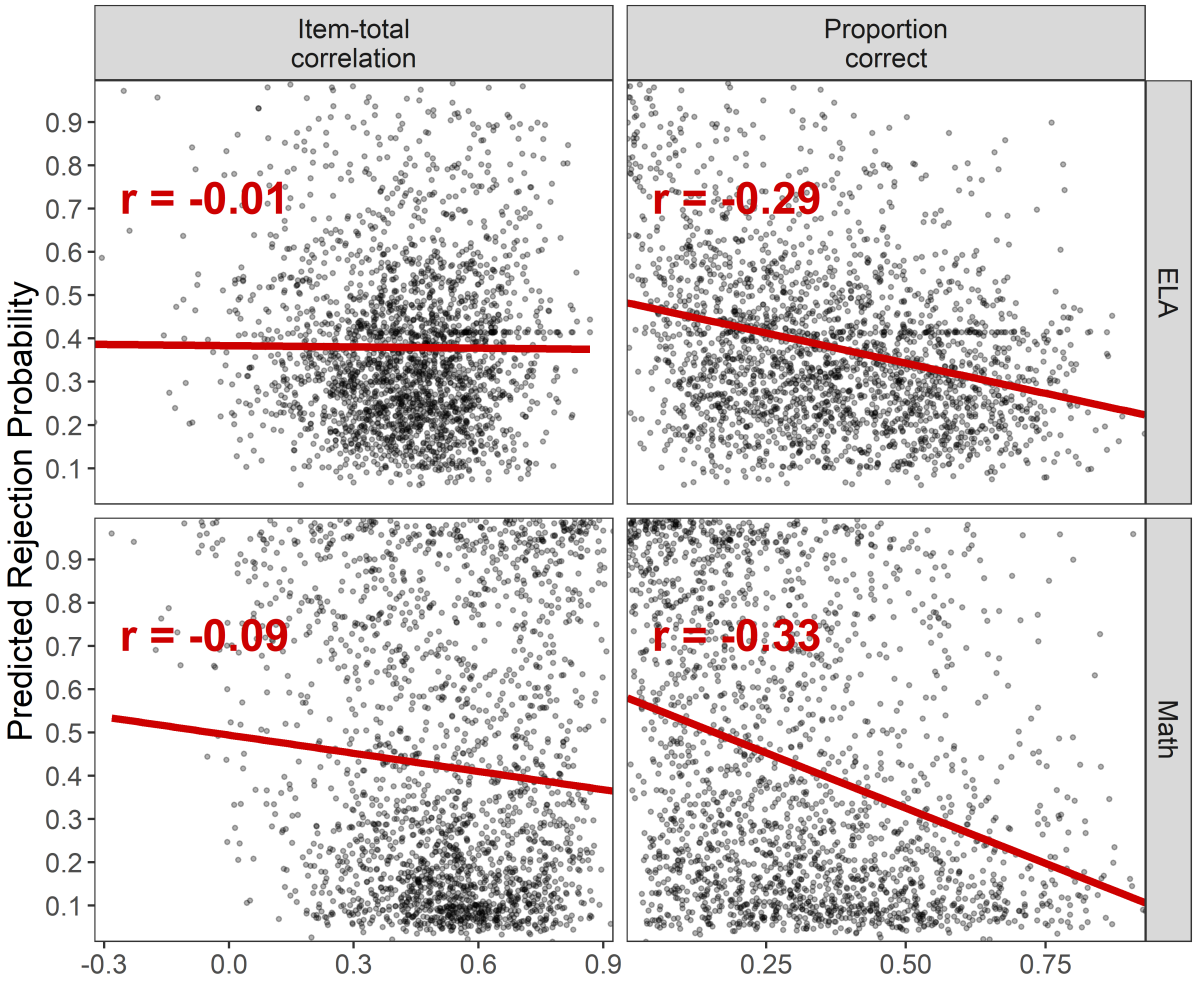
Our model takes a nearly comprehensive approach to AIE, as the key outcome of judging item quality is whether an item is acceptable for operational use. However, we cannot fully claim comprehensiveness. Item quality partly depends on the other items already in the bank. We do not

Figure 3: Boxplots of predicted rejection probability by rejected reason for items truly labeled as rejected.



Note. ELA = English language arts. S=Sensitivity (using $\geq$ .5 probability threshold), N=number of items, Scoring = issues with rubrics, answer keys, scoring logic, or scoring metadata, Abandoned = items without a suitable exam placement, Passage = rejection driven by the passage set rather than the item itself, Psychometric = psychometric issues like low discrimination and extreme difficulty, Incomplete = items that never reached field-testing due to stalling in development, Data review = questionable psychometric qualities requiring content expert review, including C-level differential item functioning, Content = content errors like standard misalignment, Bias = content accessibility, bias, fairness, or sensitivity concerns, Non-content = issues like missing, incompatible, or corrupt data that is not expected to be detectable by text-based models, (no data) = items rejected without comments or with unclear reasons. Only items truly labeled as rejected are shown. For full rejection reason details, see Appendix A.

Figure 4: Item difficulty and discrimination by fusion model-predicted rejection probability.



Note. r = Spearman correlations.

want items that are too similar to each other, which would require a similarity analysis (Peng, 2020), and the distribution of item difficulty within a bank should typically be balanced (van der Linden and Pashley, 2000). Both require comparing new items to the current bank. Some rejection reasons are also inherently unpredictable from item content alone, including item exposure, rejection due to other items within the same passage set, data integrity issues, or dependence on future events. Incorporating full item metadata such as passage text and scoring rubric could partly resolve this, but some issues will always depend on unpredictable future events. Finally, our fusion model struggled to identify items flagged for bias, sensitivity, fairness, or accessibility concerns, possibly because these judgments require contextual and cultural knowledge difficult to encode from item text alone. This suggests human review remains essential for detecting these issues in particular.

Unlike prior work focused on predicting individual item properties such as difficulty, our approach targets the operational decision that ultimately determines an item's fate, offering a single, near-comprehensive quality metric applicable to both human-written and AI-generated items. As AIG continues to expand the volume of items requiring review, scalable AIE methods like the one presented here offer practical promise for reducing the cost and burden of item evaluation in large-scale assessment programs.

References

AERA, APA, and NCME. 2014. *Standards for Educational and Psychological Testing*. American Educational Research Association, Washington, DC.

Samah AlKhuzaey, Floriana Grasso, Terry R Payne, and Valentina Tamma. 2024. Text-based question difficulty prediction: A systematic review of automatic approaches. *International Journal of Artificial Intelligence in Education*, 34(3):862–914.

Mohammad Amini, Babak Ahmadi, Xiaomeng Xiong, Yilin Zhang, and Christopher Qiao. 2025. Prompting strategies for language model-based item generation in k-12 education: Bridging the gap between small and large language models. *Preprint*, arXiv:2508.20217.

Anthropic. 2025. Claude Sonnet 4.6. `https://www.anthropic.com`. Large language model.

Luca Benedetto, Andrea Cappelli, Roberto Turrin, and Paolo Cremonesi. 2020. R2de: a nlp approach to estimating irt parameters of newly generated questions. In *Proceedings of the tenth international conference on learning analytics & knowledge*, pages 412–421.

Luca Benedetto, Paolo Cremonesi, Andrew Caines, Paula Buttery, Andrea Cappelli, Andrea Giussani, and Roberto Turrin. 2023. A survey on recent approaches to question difficulty estimation from text. *ACM Computing Surveys*, 55(9):1–37.

A. Birnbaum. 1968. Some latent trait models and their use in inferring an examinee's ability. In F. M. Lord and M. R. Novick, editors, *Statistical theories of mental test scores*, pages 397–479. Addison-Wesley.

Ruhan Circi, Juanita Hicks, and Emmanuel Sikali. 2023. Automatic item generation: foundations and machine learning-based approaches for assessments. *Frontiers in Education*, Volume 8 - 2023.

Jacob Devlin, Ming Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *NAACL HLT 2019 - 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies - Proceedings of the Conference*, volume 1.

Filipe Falcão, Patrício Costa, and José M Pêgo. 2022. Feasibility assurance: a review of automatic item generation in medical assessment. *Advances in Health Sciences Education*, 27(2):405–425.

Wanyong Feng, Peter Tran, Stephen Sireci, and Andrew Lan. 2025. Reasoning and sampling-augmented mcq difficulty prediction via LLMs. *arXiv preprint arXiv:2503.08551*.

Yanbin Fu, Hong Jiao, Tianyi Zhou, Nan Zhang, Ming Li, Qingshu Xu, Sydney Peters, and Robert W. Lissitz. 2025. Text-based approaches to item alignment to content standards in large-scale reading & writing tests.

Guher Gorgun and Okan Bulut. 2025. Instruction-tuned large-language models for quality control in automatic item generation: A feasibility study. *Educational Measurement: Issues and Practice*, 44(1):96–107.

Thomas M. Haladyna and Steven M. Downing. 1989. Taxonomy of multiple-choice item-writing rules. *Applied Measurement in Education*, 2(1):37–50.

Suhwa Han, Frank Rijmen, Allison Ames Boykin, and Susan Lottridge. 2025. Leveraging fine-tuned large language models in item parameter prediction. In *Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con): Full Papers*, pages 250–264.

Jochen Hartmann, Mark Heitmann, Christian Siebert, and Christina Schamp. 2023. More than a feeling: Accuracy and application of sentiment analysis. *International Journal of Research in Marketing*, 40(1):75–87.

Pengcheng He, Jianfeng Gao, and Weizhu Chen. 2021. Debertav3: Improving deberta using electra-style pretraining with gradient-disentangled embedding sharing. *arXiv preprint arXiv:2111.09543*.

Mads Henrichsen and Rasmus Krebs. 2025. Two-stage reasoning-infused learning: Improving classification with LLM-generated reasoning. *arXiv preprint arXiv:2507.00214*.

Cheng-Yu Hsieh, Chun-Liang Li, Chih-Kuan Yeh, Hootan Nakhost, Yasuhisa Fujii, Alexander Ratner, Ranjay Krishna, Chen-Yu Lee, and Tomas Pfister. 2023. Distilling step-by-step! outperforming larger language models with less training data and smaller model sizes. *arXiv preprint arXiv:2305.02301*.

Fu-Yuan Hsu, Hahn-Ming Lee, Tao-Hsing Chang, and Yao-Ting Sung. 2018. Automated estimation of item difficulty for multiple-choice tests: An application of word embedding techniques. *Information Processing & Management*, 54(6):969–984.

Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3.

Diana Kolesnikova, Kirill Fedyanin, Abe D Hofman, Matthieu JS Brinkhuis, and Maria Bolsinova. 2026. Estimating item difficulty with large language models as experts. *arXiv preprint arXiv:2605.18562*.

Bojie Li. 2026. Incompressible knowledge probes: Estimating black-box LLM parameter counts via factual capacity. *arXiv preprint arXiv:2604.24827*.

Ming Li, Hong Jiao, Tianyi Zhou, Nan Zhang, Sydney Peters, and Robert W Lissitz. 2025. Item difficulty modeling using fine-tuned small and large language models. *Educational and Psychological Measurement*, 85(6):1065–1090.

Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.

Zhenyan Lu, Xiang Li, Dongqi Cai, Rongjie Yi, Fangming Liu, Xiwen Zhang, Nicholas D. Lane, and Mengwei Xu. 2025. Small language models: Survey, measurements, and insights. *arXiv preprint arXiv:2409.15790*.

Ye Ma. 2025. Analyzing examinee comments using DistilBERT and machine learning to ensure quality control in exam content. *arXiv preprint arXiv:2504.06465*.

Hotaka Maeda. 2025. Field-testing multiple-choice questions with AI examinees: English grammar items. *Educational and Psychological Measurement*, 85(2):221–244.

Hotaka Maeda and Yikai Lu. 2025. Finding words associated with DIF: Predicting differential item functioning using LLMs and explainable AI. *Journal of Educational Measurement*, 62(4):883–906.

Hotaka Maeda and Yikai Lu. 2026. Multimodal test item parameter prediction from text, images, and metadata: Fusing together AI vision and language models. *Educational and Psychological Measurement*.

Eiji Muraki. 1992. A generalized partial credit model: Application of an em algorithm. *ETS Research Report Series*, 1992(1):i–30.

OpenAI. 2024. GPT-4o system card. *arXiv preprint arXiv:2410.21276*.

Radek Pelánek, Tomáš Effenberger, Adam Kukučka, et al. 2022. Towards design-loop adaptivity: identifying items for revision. *Journal of Educational Data Mining*, 14(3):1–25.

Fang Peng. 2020. *Automatic enemy item detection using natural language processing*. Ph.D. thesis, University of Illinois Chicago.

Jim Prentzas and Ariadni Binopoulou. 2025. Explainable artificial intelligence approaches in primary education: A review. *Electronics*, 14(11).

Qwen Team. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.

Alexander Scarlatos, Nigel Fernandez, Christopher Ormerod, Susan Lottridge, and Andrew Lan. 2025. Smart: Simulated students aligned with item response theory for question difficulty prediction. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 25082–25105.

Jinnie Shin and Mark J Gierl. 2024. Automated short-response scoring for automated item generation in science assessments. In *The Routledge International Handbook of Automated Essay Evaluation*, pages 504–534. Routledge.

Bin Tan, Nour Armoush, Elisabetta Mazzullo, Okan Bulut, and Mark Gierl. 2025. A review of automatic item generation techniques leveraging large language models. *International Journal of Assessment Tools in Education*, 12(2):317–340.

Wim J. van der Linden and Peter J. Pashley. 2000. *Item Selection and Ability Estimation in Adaptive Testing*, pages 1–25. Springer Netherlands, Dordrecht.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *arXiv preprint arXiv:1706.03762*.

Yu Wang, Madhumitha Gopalakrishnan, and Yoav Bergner. 2025. Using generated rubrics to provide a window into item evaluation with multi-agent LLMs. In *International Conference on Artificial Intelligence in Education*, pages 203–217. Springer.

Benjamin Warner, Antoine Chaffin, Benjamin Clavié, Orion Weller, Oskar Hallström, Said Taghadouini, Alexis Gallagher, Raja Biswas, Faisal Ladhak, Tom Aarsen, Nathan Cooper, Griffin Adams, Jeremy Howard, and Iacopo Poli. 2024. Smarter, better, faster, longer: A modern bidirectional encoder for fast, memory efficient, and long context finetuning and inference. *arXiv preprint arXiv:2412.13663*.

Kelly Wauters, Piet Desmet, and Wim Van Den Noortgate. 2012. Item difficulty estimation: An auspicious collaboration between data and judgment. *Computers & Education*, 58(4):1183–1193.

Victoria Yaneva, Le An Ha, Peter Baldwin, and Janet Mee. 2020. Predicting item survival for multiple choice questions in a high-stakes medical exam. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 6812–6818, Marseille, France. European Language Resources Association.

Martin C Yu and Maura I Burke. 2026. Automatic item generation, evaluation, and scale construction of non-cognitive measures with generative language models. In Isaac Thompson, Georgi P Yankov, and Ivan Hernandez, editors, *Artificial Intelligence for I-O Psychologists: Research and Applications*. Oxford University Press.

Appendix A Rejection Reason Definitions

1. Content: content issue or error, including but not limited to (1) misalignment with the blueprint or standards, (2) unrealistic, inconsistent, or confusing scenarios, or (3) basic issues with the text or image such as grammar errors or missing information. A large portion of these items are rejected at a content review event just prior to field testing.

2. Psychometric: The following statistical criteria result in immediate rejection without manual review: low proportion correct, low item-total correlation, a distractor with a higher item-total correlation than the correct answer, poor inter-rater reliability on hand-scored items, or polytomous items where mean estimated θ does not increase monotonically across score levels. Items can also be rejected for less severe psychometric issues, though these require review and situational judgment: low or negative discrimination, too easy, too difficulty, insufficient sample size, overexposure, parameter drift, or differential item functioning.

3. Data review: These are items with questionable psychometric qualities that needed to be reviewed by content experts. Items flagged with C-level differential item functioning are always reviewed by content experts.

4. Bias: Content bias, sensitivity, fairness, or accessibility reasons

5. Abandoned: Items that could not find a suitable use in an exam. Many tend to be extra items in passage sets. Sometimes, the item itself has no content issues.

6. Incomplete: Items stuck in the item development process. Never reached field-testing.

7. Passage: The entire passage set (i.e., stimulus or testlet) was rejected. Sometimes, the item itself has no content issues.

8. Scoring: Any issue related to scoring rubric, answer key, scoring logic, or scoring meta data. This can be a combination of content, data integrity, and psychometric issues.

9. Non-content: Issues unrelated to the text or images of the item itself, such as missing metadata, corrupt or malformed data, format incompatibilities, import errors, or expired copyright permissions. Because these are primarily data integrity issues, text-based models are not expected to detect them.

10. (no data): Items rejected without comments or unclear reasons.

Appendix B Qwen3 Zero-Shot Classification Prompt

Below is an example python script for generating the Qwen3-0.6B zero-shot classification prompt.

```
SYSTEM_PROMPT = (
    "You are an expert assessment item reviewer. You will be given an assessment item. Respond with
    exactly one word: 'accept' or 'reject'. Respond 'accept' if the item is usable for assessment as-
    is. Respond 'reject' if the item has issues that go beyond a simple fix, such as: fundamental
    ambiguity, significant bias or sensitivity concerns, clear misalignment to the stated standard,
    or broken/unfixable formatting. Do not include any explanation, punctuation, or additional text.
    "
)
def make_messages(row):
    return {
        "messages": [
            {"role": "system", "content": SYSTEM_PROMPT},
            {"role": "user", "content": f"This is a {row['subject']} subject, grade {row['grade']}
    item: \n{row['text']} \n{row['standards']} \nAny relevant images, figures, or tables have been
    excluded."
                }]}
```

Appendix C Qwen3 Item Critique Prompt

Below is an example python script for generating the Qwen3-0.6B item critique prompt.

```
SYSTEM_PROMPT = (
    "You are an expert assessment item reviewer with deep knowledge of psychometrics and item
    development best practices. You will be given an assessment item. Review the item for common
    quality issues such as: ambiguity, item difficulty, cultural or demographic bias and sensitivity,
     content alignment to standards, and flawed grammar or formatting. Write a concise 2 sentence
    summary: if the item has issues, name them plainly and briefly explain why they matter; if the
    item looks sound, say so and note its strongest quality. Do not use bullet points or headers. Be
     direct and specific - avoid vague praise or vague criticism."
)
def make_messages(row):
  return {
    "messages": [
            {"role": "system", "content": SYSTEM_PROMPT},
            {"role": "user", "content": f"This is a {row['subject']} subject, grade {row['grade']}
    item: \n{row['text']} \n{row['standards']}  \nAny relevant images and figures or tables have
    been excluded. \nWrite a concise 2 sentence summary about the item quality, including but not
    limited to its item difficulty, bias and sensitivity, and content alignment to standards."
            }]}
```

Appendix D Example Qwen3 Critiques

Table 2: Example Critiques Generated using Qwen3-0.6B.

Subject	Status	Grade	Critique	Sentiment	P_{fusion}	$P\Delta$
ELA	Accept	3	The item presents a coherent analysis of Samantha's emotional state, aligning well with grade 3 standards. It includes specific details about her actions and observations, making it suitable for assessing character responses to a narrative. No major quality issues are identified, as the content directly supports the intended learning outcomes.	6.8	.26	-0.03
ELA	Accept	3	The item contains a grammatical issue with repeated use of "sentence" and ambiguous phrasing, making clarity challenging. It aligns well with standards for grade 3 ELA, particularly focusing on analyzing stories and identifying central messages through textual evidence. The content is appropriate and includes elements like images and tables as specified.	6.8	.66	0.34
ELA	Reject	8	The item contains repetition (e.g., "very tangled" vs. "tangled") and lacks variation in language precision, making it a flawed grammar issue. It aligns with standard reading comprehension tasks but fails to address ambiguity in the task instructions.	-7.5	.73	0.39
ELA	Reject	11	The item presents a clear connection between Eisenhower's plan for America's development and his focus on global responsibility, aligning with the standard of analyzing how ideas interact over time. It correctly identifies key points such as the paragraph's purpose to emphasize mutual trust and respect, making it a sound assessment item.	6.8	.37	0.02
Math	Accept	5	The item contains ambiguity in the number "about 90," which may leave students unsure if it refers to the total number of people or a specific group. It also shows potential difficulty in the answer key, requiring clarification on what data is needed to estimate waiting time. The question aligns with grade-level standards for mathematical estimation skills.	-4.9	.86	0.36
Math	Accept	6	The item lacks necessary visual elements (no image/table) and does not provide the required data (number of neighbors), making it ambiguous. It aligns with standards like reporting observations but fails to include supporting context.	-7.5	.14	-0.05
Math	Reject	6	The item presents a clear unit rate question about distance traveled, aligning with grade-level standards. It includes relevant data (3000 feet in one hour) without ambiguity, and the context is neutral, making it suitable for a math curriculum. No significant biases or flaws were identified.	6.8	.54	-0.02
Math	Reject	11	The item presents a valid mathematical scenario where the slopes of AE and BC are consistent with the properties of a right triangle, and all stated facts align with the given conditions. It does not contain ambiguities, biases, or content misalignment with educational standards. The highest quality is the clear application of slope relationships in a geometric context.	6.8	.70	0.44

Note. Status = true status label, Sentiment = sentiment logit (positive = positive sentiment, negative = negative sentiment), P_{fusion} = fusion model probability of rejection, $P\Delta$ = P_{fusion} minus text-only model probability of rejection, ELA = English language arts. Two critiques from accepted and rejected ELA and math items were randomly selected, with the requirement that one had a near-zero $P\Delta$, and another had $P\Delta > .3$.

Appendix E Classification with .25 Cutoff

Table 3: Classification Performance by Model with .25 Cutoff Threshold

Model	p	AUC	Accuracy	Precision	Sensitivity	Specificity	F1
Raw text-only	.78	.76	.51	.41	.92	.30	.56
ELA	.84	.69	.42	.33	.94	.20	.49
Math	.71	.80	.61	.50	.91	.42	.64
Critique-only	.80	.73	.48	.39	.92	.26	.55
ELA	.92	.63	.35	.31	.96	.09	.47
Math	.67	.79	.62	.50	.88	.46	.64
Fusion (raw text + critique)	.69	.80	.58	.45	.90	.42	.60
ELA	.75	.72	.48	.35	.88	.31	.51
Math	.62	.86	.70	.57	.91	.56	.70
ELA Fusion (stand-alone)	.89	.70	.38	.32	.96	.13	.48
Math Fusion (stand-alone)	.59	.85	.71	.58	.88	.59	.70

Note. ELA = English language arts, p = proportion of items predicted as reject, AUC = area under the curve. Cutoff threshold for models were fixed to .25.

Production-Ready Automated Item Generation in Educational Assessment: Integration with Operational Workflows

Hotaka Maeda [1,2]*
Kargi Chauhan [2]
Tharunya Chandrashekar[1]

[1]Smarter Balanced, Santa Cruz, California, United States of America
[2]University of California-Santa Cruz, Santa Cruz, California, United States of America

Abstract

We present a production-ready automated item generation pipeline integrated into a large-scale assessment program's workflows. A one-shot approach prompts LLMs from an operational source item and existing guidelines, then populates metadata, screens quality, and uploads items for human review. Generated items passed expert review, and difficulty prompting reliably shifted difficulty.

1 Background

Large-scale educational assessments require hundreds or thousands of new items annually, yet traditional item development remains slow and expensive. Recent work has shown that automated item generation (AIG) using large language models (LLMs) can produce structurally plausible items across a range of subjects and formats (Circi et al., 2023; Falcão et al., 2022; Tan et al., 2025; Omopekunola et al., In press). However, most AIG research treats generation as a standalone task, producing items in isolation from the testing program that would ultimately use them. A production-ready AIG pipeline instead requires operational integration. This is challenging when the generation process must be embedded into a preexisting examination program that has run for years and already has its own item requirements, data, platforms, and human workflows. Established programs have guidelines defining what items must look like, an item bank database holding operational items, and a content team that writes and revises items through interfaces they use daily.

This challenge is not unique to assessment. Organizations broadly have struggled to translate AI adoption into realized value, a pattern often described as the AI productivity paradox (McElheran et al., 2025; Brynjolfsson et al., 2021, 2017). Evidence consistently points to organizational and process factors rather than AI model capability as the binding constraint. The challenge is acute in established organizations whose existing routines, legacy systems, and layered processes are hard to adapt (Brynjolfsson et al., 2021). Educational testing programs are a clear example of this setting.

Framing production-readiness as integration has concrete consequences for how the pipeline is built. First, because the program already maintains documented item requirements, these can be fed directly into LLMs at generation time rather than reinvented. Second, because the program already has a bank of operational items, a single existing item can serve as a one-shot source item, which lets generated items inherit the source item's metadata rather than requiring it to be reconstructed. Third, because items are stored and reviewed in an existing database, generated items can be imported back into that same system, keeping data organized and letting content experts review new items in a familiar platform. Fourth, because items are stored as code rather than rendered form, this round trip through the database is what renders them into the visual format human reviewers expect (see Figure 1). Together, these design choices let an established program adopt AIG with minimal disruption and added effort while still holding items to its existing standards.

2 Aims

We present an AIG pipeline for established large-scale standardized English Language Arts (ELA) and mathematics assessments designed for production use. Rather than treating generation as a standalone exercise, the system is designed to fit the infrastructure, policies, and people of the Smarter Balanced testing program. After introducing the pipeline, we evaluate it in two studies. Study 1 is a

*Correspondence concerning this article should be addressed to Hotaka Maeda, Smarter Balanced, University of California-Santa Cruz, 1156 High St, Santa Cruz, CA 95064. E-mail: hotaka.maeda@smarterbalanced.org

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 130–139
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

(a) Generated JSON


```
"options": [
  {
    "isCorrectAnswerKey": false,
    "text": "<p>81</p>"
  },
  {
    "isCorrectAnswerKey": true,
    "text": "<p>108</p>"
  },
  {
    "isCorrectAnswerKey": false,
    "text": "<p>112</p>"
  },
  {
    "isCorrectAnswerKey": false,
    "text": "<p>120</p>"
  }
],
"prompt": "<p>The class collects 12 cans on Monday and 15 cans on Tuesday
      . They collect the same total each week for 4 weeks.</p><p><br/>
      How many cans do they collect in 4 weeks?</p>"
```


(b) Rendered item

The class collects 12 cans on Monday and 15 cans on Tuesday. They collect the same total each week for 4 weeks.

How many cans do they collect in 4 weeks?

Ⓐ 81

Ⓑ 108

Ⓒ 112

Ⓓ 120

Figure 1: Example generated multiple choice item. The large language model generates items in HTML format embedded within a JSON format. Once uploaded to the item bank database, the rendered item can be viewed for human review on the front-end interface of the database, the same platform the content team already uses to write and revise items. Only a portion of the generated JSON is shown.

small initial check of whether the generated items are acceptable to human content experts. Study 2 uses a larger generated item pool to examine how the generation inputs affect item quality. We vary the content specifications, fairness guidelines, and target difficulty, and evaluate the resulting items with automated screening and prediction methods.

3 Proposed AIG Pipeline

The pipeline is organized into seven modular sequential stages.

1. **Data retrieval:** We fetch a manually selected operational item from a database using queries. This source item plays a role analogous to the item model in traditional automated item generation (Gierl and Lai, 2012), defining the structure, style, metadata, and covered content standard for the newly generated items. Unlike a conventional item model, which is a manually engineered template, the source item is an existing operational item used directly as a one-shot example. We also retrieve related item development documents, including the content specifications and sensitivity and fairness guidelines that govern the existing testing program.

2. **One-shot prompt construction:** All retrieved information is converted into a prompt. HTML format embedded within JSON is used, which is both well suited for LLM processing and consistent with the format used to store items in the item database. The LLM is instructed to mirror the organization and conventions of existing items by following the structured JSON format of the one-shot source item, and instructed to return its response in the same structured format (see Appendix A). Using the database's own storage format at both input and output is what allows generated items to move back into the operational system later without reformatting.

3. **Item generation:** Send the prompt to OpenAI GPT-5.6 Terra through an API (OpenAI, 2026). The LLM generates text content of multiple items alongside scoring rubrics in a structured JSON format.

4. **Metadata Assignment:** The one-shot approach allows the LLM to mimic many features of the source item, which is an advantage because the metadata of the source item can be deterministically copied to the generated items. While items can have hundreds of back-end metadata fields, examples include the subject, grade level, blueprint categories, standard alignment, or item type. This significantly reduces the burden of post-generation item data management. Note that not all metadata and auxiliary data can be copied because they would be different for the new items (e.g., item parameters or text-to-speech audio data).

5. **Natural language processing (NLP) quality control (QC):** We compare generated items against the known distribution of operational item characteristics using ten NLP features computed from item text using the `textstat` Python package (Bansal and Aggarwal, 2025). These included word count, letter count, sentence count, words per sentence, average sylla-

bles per word, average letters per word, Flesch Reading Ease, Automated Readability Index (ARI), Dale-Chall readability score, and estimated reading time. Specifically, we confirm that these variables are within the minimum and maximum range observed in operational items within the same subject, grade, item type, and claim. Items that do not meet the required ranges are rejected outright.

6. **AI review:** Similar prompt from Step 2 where the source items are replaced with the generated items are sent to GPT-5.6 Terra again to re-check the fidelity to all requirements (see Appendix B). The LLM classified each item as "accept", "flag", or "reject", and generated a short rationale. The returned responses are provided to the content experts to help assist during the human review phase.

7. **Item bank database upload:** Upload accepted items to the secure item database environment along with the metadata (i.e., the same database that stores the source item) through its API. Successful upload serves an additional quality assurance check, as the data format must match what is expected by the database. Critically, items are stored as code rather than in rendered form, so importing them back into the database is what renders them into the visual format reviewers expect. Uploaded items are available for immediate render and human review on the front-end interface of the database, the same platform the content team already uses to write and revise items.

4 Study 1: Initial Acceptability

An initial small study was conducted to ensure that the AIG system was functioning acceptably well prior to a more in-depth evaluation. Four different types of source items were used for the study: (1) ELA multiple-select (MS) and (2) hot-text-question-select (HTQS; selection of words within paragraphs), and (3) math multiple-choice (MC) and (4) equations. The four source items were used to generate ten items each ($4 \times 10 = 40$ items) and were sent to the ELA and math content experts for human review.

5 Study 1: Results

A total of 8/10 ELA multiple-select items, 10/10 hot-text-question-select items, 6/10 Math multiple-choice, and 10/10 equation items passed NLP QC. ELA and math content experts reviewed five items for each of the item types ($4 \times 5 = 20$ items). They all passed human review by three reviewers per subject, and were therefore all usable for operations with little to no modification. However, a key challenge was that flaws in source items were sometimes inherited by the generated items. This makes careful source item selection essential, as well as revision of any flawed source items before generation begins. LLM usage costs were less than 2 cents per generated item.

6 Study 2: Ablations and Item Difficulty

Study 1 demonstrated that the AIG pipeline can produce operationally viable items integrated into existing workflows. Study 2 examines how the two primary guideline inputs, content specifications and fairness and sensitivity guidelines, along with difficulty prompting, affect the quality and linguistic characteristics of generated items. We generated a larger item pool under systematically varied conditions, characterized the items using NLP, and evaluated their operational viability using prediction models.

6.1 Item Generation Design

We used a fully crossed experimental design with three factors. The first factor was the inclusion of content specifications, either present or absent. The second factor was the inclusion of the fairness and sensitivity guidelines, either present or absent. The third factor was the target difficulty, where the prompt specified to generate items easier, similar, or harder difficulty compared to the source item. Crossing these factors produced $2 \times 2 \times 3 = 12$ conditions. When one or both guideline inputs were absent, the corresponding portion of the prompt was removed entirely.

Ten operational source items served as one-shot anchors, five for ELA (all grade 3, 3 HTQS, 1 MC, 1 MS item) and five for mathematics (all grade 4, 3 MC, 2 equations). Each source item was used to generate items under all 12 conditions, and 10 items were generated per source item within each condition. This yielded $12 \times 10 \times 10 = 1{,}200$ generated items in total, balanced across the two subjects. All items were generated using the pro-

posed pipeline, with the difficulty instruction and the guideline ablations applied at the prompt construction stage.

6.2 Three Automated Screening Criteria

We characterized the linguistic patterns of generated items using ten NLP features computed from item text in the NLP QC step. As described in the proposed pipeline, items were labeled as "reject" when any of these variables were outside of the minimum and maximum range observed in operational items within the same subject, grade, item type, and claim. In addition, AI reviewer results are analyzed by treating "accept" and "flag" items as "accept", and rejecting the "reject" items.

To assess the operational viability of the generated items, we applied the automated item evaluation (AIE) fusion model (Maeda and Lu, 2026a) to predict whether an item would be accepted or rejected for operational use. It was trained on 52,759 ELA and mathematics items from the same testing program targeted in the current paper. The fusion model combines two fine-tuned DeBERTaV3-large encoders (He et al., 2021). The first encoder processes the raw item text. The second encoder processes a two-sentence item quality critique generated by a Qwen3-0.6B language model (Qwen Team, 2025). The 1,024-dimensional final-layer representations from both encoders are concatenated and passed through a fusion head to produce a rejection probability. On held-out test data, predictions achieved stronger performance for math (AUC = .86) than for ELA (AUC = .72).

The AIE model was used to produce a rejection probability for all 1,200 generated items. Probabilities at or above .25 (rather than .50) were treated as rejected. This low threshold has a high sensitivity (.88 for ELA and .91 for math) with reduced specificity (.31 for ELA and .56 for math), which is recommended in an AIG context where generating additional items has a low cost (Maeda and Lu, 2026a).

6.3 Predicting Item Discrimination and Difficulty

To assess whether difficulty prompting shifted item difficulty properly, and whether the automated screening methods reject items with low discrimination, we predicted item discrimination (a) and difficulty (b) using a prediction model (Maeda and Lu, 2026b). We use the text-only variant of the model (i.e., no images and no item-level metadata input), which was a fine-tuned DeBERTaV3-large (He et al., 2021) that converts each item text component into a 128-dimensional representation, aggregates these into an item-level representation using attention pooling, and predicts item response theory item parameters through a regression head. Trained on the same standardized test data targeted in the current paper, predicted discrimination and difficulty achieved R^2 values of .45 and .38, respectively.

6.4 Statistical Analysis

We tested whether the experimental conditions affected the generated items using mixed-effects models fit in R with the `lme4` package (Bates et al., 2015). Each generated item was nested within the source item it was derived from, so all models included a random intercept for source item to account for the non-independence of items sharing an anchor. The three experimental factors were entered as fixed effects. Continuous outcomes, including the NLP metrics, predicted item difficulty, and predicted discrimination, were analyzed with linear mixed models. The three binary screening criteria (NLP, AI reviewer, and AIE model) were analyzed with a generalized linear mixed model using a logit link and a binomial response distribution. All models were estimated by maximum likelihood so that models differing in their fixed-effects structure could be compared directly.

To determine whether the experimental factors affected the outcomes, using likelihood ratio tests, we compared the (1) full model with all main effects and 3-way interactions, (2) reduced model with only main effects, and (3) a null model with only the intercept. Fixed-effect significance ($\alpha = .05$) was assessed using Satterthwaite-approximated degrees of freedom for the linear models via the `lmerTest` package (Kuznetsova et al., 2017). NLP metrics were compared only across the difficulty condition. Spearman correlations between the screening criteria and predicted item discrimination are also reported. When calculating correlations with AI reviewer decisions, all three categories were used (accept=0, flag=1, reject=2).

7 Study 2: Results

7.1 Screening Evaluation

We applied three independent screening criteria to the generated items. With a mean predicted rejec-

tion probability of .35 ($SD = .20$), 75% of ELA and 50% of math items were rejected by the AIE model. These figures are only 6 percentage points above the model's false positive rate, which provides partial evidence that these items are generally good quality. The AI reviewer rejected 46% of items, and NLP QC rejected 16%. Only 13% of ELA and 22% of math items passed all three criteria. NLP and AI review were modestly related ($\rho = .19$), while the AIE model was essentially uncorrelated with both NLP ($\rho = .06$) and AI review ($\rho = -.05$). These low correlations show that the three screening tools are not redundant and each may catch distinct item issues. The AI reviewer ($\rho = -.44$), NLP ($\rho = -.20$), and AIE ($\rho = -.11$) screening criteria all tended to reject items that had low predicted item discrimination.

Interactions effects were not present for the NLP method, $\chi^2(7) = 2.7$, $p = .91$, nor AI reviewer, $\chi^2(7) = 1.91$, $p = .96$, while all three main effects were significant for both ($p < .05$) except for the bias guideline effects for the NLP method ($p = .75$). Predicted AIE rejection rates did not differ across any experimental condition, $\chi^2(11) = 4.5$, $p = .95$.

Rejection rates differed markedly by subject. The AI reviewer rejected math items (61%) far more often than ELA items (32%). The NLP-based QC showed the reverse, rejecting more ELA items (21%) than math items (11%). For ELA, supplying the content specifications decreased AI reviewer rejection rates from 38% to 25%, and NLP rejection rate from 24% to 17%. AI reviewer rejected more ELA items when items were prompted to be easier (40%) than similar (31%) or harder (24%). The bias guidelines slightly raised the ELA rejection rate (37% versus 26%) rather than lowering it. For math, AI reviewer and NLP rejection rates were similar across all conditions. The only exception was that the NLP method tended to flag harder (14%) and similar (14%) math items than easier ones (6%).

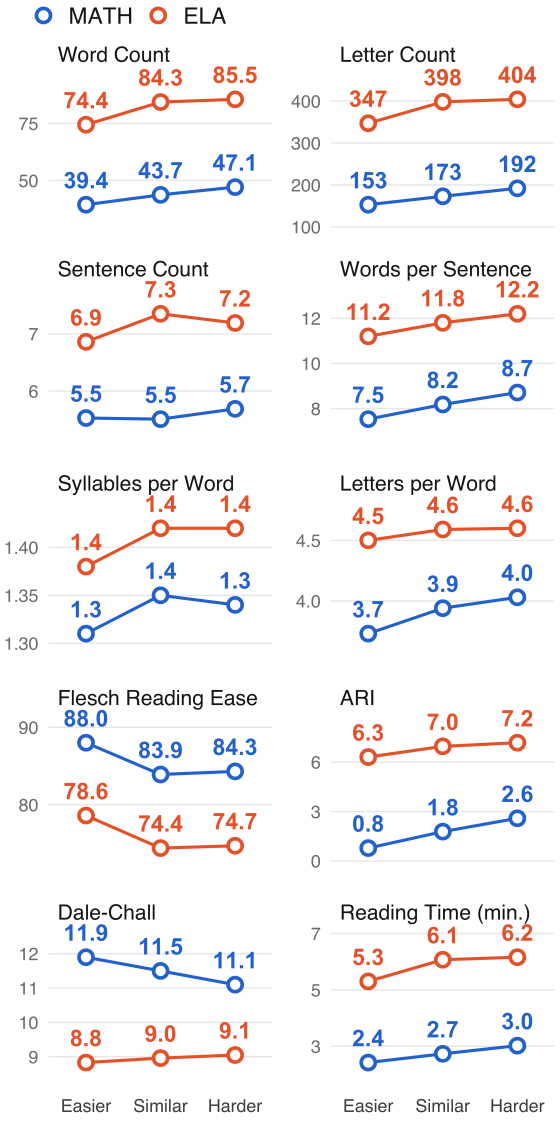


Figure 2: NLP metrics by difficulty prompting condition. NLP = natural language processing, ELA = English language arts, ARI = Automated Readability Index. Prompt specified to generate items easier, similar, or harder difficulty compared to the source item.

7.2 NLP Metrics

All ten NLP metrics were statistically significantly affected by the difficulty conditions ($p < .05$). The generated items' linguistic features shifted with prompted difficulty, providing convergent evidence that the difficulty manipulation operated as intended (see Figure 2). For mathematics, mean word count rose from 39 to 44 to 47 words across the easier, similar, and harder conditions, while for ELA it rose from 75 to 84 to 86 words. The ARI showed a corresponding rise from 0.8 to 2.6 for mathematics and from 6.4 to 7.2 for ELA. Letter count, words per sentence, and Flesch Reading Ease all moved in the expected direction for both subjects as well, indicating that the difficulty effect reflects a genuine change in item complexity rather than an artifact of any one measure. For several metrics the largest step occurred between the easy and medium conditions, with little further change from medium to hard. The only unexpected result was the Dale-Chall score for mathematics, which indicated simpler vocabulary in harder items. However, Dale-Chall is designed to quantify reading

comprehension difficulty, and may not be effective for gauging math item difficulty.

7.3 Predicted Item Discrimination and Difficulty

Interactions effects of the conditions were not present for item discrimination, $\chi^2(7) = 9.0$, $p = .25$, nor difficulty, $\chi^2(7) = 8.04$, $p = .33$, while all three main effects were significant for both ($p < .01$). Items had 0.06 higher predicted discrimination when prompted for easier difficulty compared to harder. Difficulty prompting shifted predicted item difficulty in the intended direction (see Figure 3). Items prompted to be easier, similar difficulty, or harder than the source items averaged -0.85, -0.55, and 0.20 for math, and -0.99, -0.72, and -0.69 for ELA, respectively (see Table 1). Difficulty prompting alone accounted for 10.0% of the variance in predicted difficulty.

Generated items prompted to be similar difficulty, $M(SD) = -0.75(0.50)$, tended to be more difficult than their source items $-1.07(0.42)$. Prompting for easier items, $-1.01(0.40)$, produced the closest match to source item difficulty. Items became modestly less discriminating and more difficult when content specifications and bias guidelines were supplied ($p < .01$). This is consistent with the expectation that adding constraints to the generation prompt can trade off against other item qualities.

8 Discussion

This study presents a successful end-to-end AIG pipeline designed for production use that adapts to ELA and mathematics configurations within the Smarter Balanced assessment program. The key design decision is that our AIG system takes advantage of the item development structures that were already in place. The existing item bank platform, source items, content specifications, and sensitivity and fairness guidelines provide a backbone to generating operationally viable items. We also used prediction models (Maeda and Lu, 2026a,b) trained on historical assessment program data to evaluate the generated items. Using these models, NLP, and AI review, newly generated items are automatically vetted. Items with high chance of being good quality are uploaded to the item bank, where they are rendered and immediately available for human review. The result is a smooth transition from machine generation to human review. Our system is

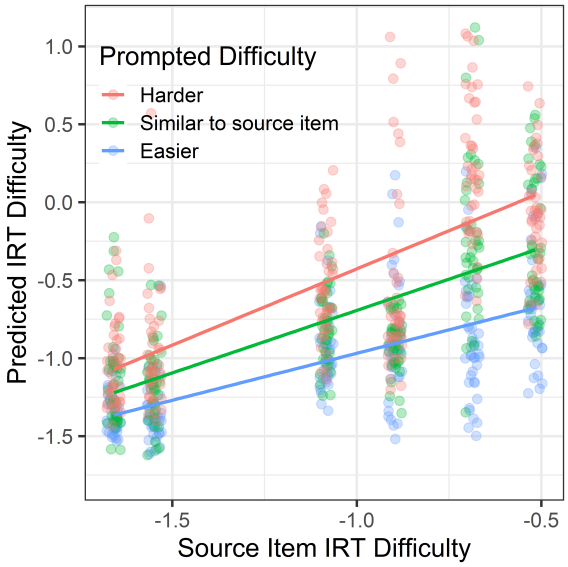


Figure 3: Predicted difficulty of generated items relative to the source item difficulty and prompt. IRT = item response theory. Predicted IRT b-parameters based on the 2-parameter logistic model are from a text-based prediction model (Maeda and Lu, 2026b). Shown are 721 items generated from 6 of the 10 source items that had item parameters.

also highly adaptable and modular. For example, we generated four different item types, while many AIG systems are designed to generate only multiple choice items (Gierl et al., 2012; Song et al., 2025). Our system is item-type agnostic, relying only on the source item and content specification. This design lowers the friction of transitioning an assessment program to adopt an AIG system.

The three automated screening criteria all tended to reject items with poor predicted item discrimination, but the criteria were not redundant. Each appears to capture a distinct class of item issues. The NLP-based QC method catches objective deviations from the operational norm, such as excessive length. The AI reviewer was more strongly associated with predicted item discrimination, suggesting it may identify items that are not designed cohesively, such as response options that do not function well together. The AIE model was trained to reproduce the testing organization's historical accept and reject decisions, so it arguably rejects items for the broadest range of reasons among the three. Because the criteria are complementary rather than overlapping, all three could be incorporated into the QC steps of the AIG pipeline.

Difficulty prompting effectively controlled the predicted difficulty of the generated items. Change

		ELA		Math	
Factor	Category	Discrimination	Difficulty	Discrimination	Difficulty
Content specs	Absent	0.67 (0.10)	−0.87 (0.63)	0.72 (0.17)	−0.56 (0.54)
	Present	0.63 (0.12)	−0.73 (0.65)	0.72 (0.18)	−0.48 (0.54)
Sensitivity and Fairness Guide	Absent	0.66 (0.10)	−0.84 (0.60)	0.73 (0.18)	−0.55 (0.56)
	Present	0.64 (0.13)	−0.76 (0.68)	0.72 (0.17)	−0.50 (0.51)
Difficulty	Easier	0.70 (0.11)	−0.99 (0.61)	0.74 (0.15)	−0.85 (0.38)
	Similar	0.64 (0.12)	−0.72 (0.65)	0.73 (0.17)	−0.53 (0.49)
	Harder	0.62 (0.11)	−0.69 (0.63)	0.70 (0.20)	−0.20 (0.53)

Table 1: Predicted discrimination and difficulty by condition and subject. ELA = English language arts. Values are M (SD) of 2-parameter logistic model parameters, predicted using a model introduced by Maeda and Lu (2026b). Higher b indicates a more difficult item. Main effects of each factor was statistically significant ($p < .01$).

was also evident in all linguistic features of the items, which provide convergent evidence that the difficulty manipulation worked as intended. One consistent departure was that generated items tended to be slightly more difficult than the source item even when prompted for similar difficulty.

Supplying content specifications lowered ELA rejection rates, indicating that faithfulness of item structure can be reviewed reasonably well with a machine-based method. However, whether the fairness guidelines improved the items was less clear. Their effect was difficult to assess with any of the three automated screening criteria, and including them did not meaningfully change screening rates. This leaves two possibilities. Either the guidelines did not substantially change the generated items, or the screening methods were unable to detect the changes they produced. The latter is plausible given that the AIE model detects bias, sensitivity, and fairness concerns poorly (Maeda and Lu, 2026a). Our field consistently warns that we cannot place blind trust in AI-generated content in educational measurement (Bulut et al., 2024; Yin et al., 2025). Human review should therefore follow every AIG process, with reviewers concentrating their attention on the fairness, bias, and sensitivity issues that machines handle least reliably. All items from Study 1 passed human review, providing some evidence that our system is typically able to avoid these concerns.

A key limitation of our pipeline is the inability to generate complex items such as those involving images, graphs, tables, or testlets, and the reliance on high-quality preexisting source items, which prevents the generation of entirely new item types. Given the modular design of the pipeline, several of these shortcomings could be addressed through further development. Another limitation was that only 6 of the 10 source items (4 mathematics and 2 ELA) had operational item parameter data available for analysis. Finally, the difficulty and discrimination comparisons depend on predicted item parameters, which should be taken with caution.

Acknowledgments

We gratefully acknowledge the contributions of Alec Gross, Alex Dean, Ana Stanciu, Aren Simmons, Brad Vasil, Catherine Francis, Danielle Ricks, Gilbert Salazar, Katie Westgate, Natasha Johnson, Shumin Jing, and Valerie Pannell.

References

Shivam Bansal and Chaitanya Aggarwal. 2025. textstat.

Douglas Bates, Martin Mächler, Ben Bolker, and Steve Walker. 2015. Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1):1–48.

Erik Brynjolfsson, Daniel Rock, and Chad Syverson. 2017. Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics. NBER Working Paper 24001, National Bureau of Economic Research, Cambridge, MA.

Erik Brynjolfsson, Daniel Rock, and Chad Syverson. 2021. The productivity J-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics*, 13(1):333–372.

Okan Bulut, Maggie Beiting-Parrish, Jodi M Casabianca, Sharon C Slater, Hong Jiao, Dan Song, Christopher M Ormerod, Deborah Gbemisola Fabiyi, Rodica Ivan, Cole Walsh, et al. 2024. The rise of artificial intelligence in educational measurement: Opportunities and ethical challenges. *Chinese/English Journal of Educational Measurement and Evaluation*, 5(3).

Ruhan Circi, Juanita Hicks, and Emmanuel Sikali. 2023. Automatic item generation: foundations and machine learning-based approaches for assessments. *Frontiers in Education*, 8.

Filipe Falcão, Patrício Costa, and José Miguel Pêgo. 2022. Feasibility assurance: a review of automatic item generation in medical assessment. *Advances in Health Sciences Education*, 27(2):405–425.

Mark J. Gierl and Hollis Lai. 2012. The role of item models in automatic item generation. *International Journal of Testing*, 12(3):273–298.

Mark J Gierl, Hollis Lai, and Simon R Turner. 2012. Using automatic item generation to create multiple-choice test items. *Medical education*, 46(8):757–765.

Pengcheng He, Jianfeng Gao, and Weizhu Chen. 2021. Debertav3: Improving deberta using electra-style pretraining with gradient-disentangled embedding sharing. *arXiv preprint arXiv:2111.09543*.

Alexandra Kuznetsova, Per Bruun Brockhoff, and Rune Haubo Bojesen Christensen. 2017. lmerTest package: Tests in linear mixed effects models. *Journal of Statistical Software*, 82(13):1–26.

Hotaka Maeda and Yikai Lu. 2026a. Automated item evaluation: Predicting item acceptance and rejection using LLM-generated critiques. *arXiv preprint arXiv:2608.06609*.

Hotaka Maeda and Yikai Lu. 2026b. Multimodal test item parameter prediction from text, images, and metadata: Fusing together AI vision and language models. *Educational and Psychological Measurement*.

Kristina McElheran, Mu-Jeung Yang, Zachary Kroff, and Erik Brynjolfsson. 2025. *The Rise of Industrial AI in America: Microfoundations of the Productivity J-curve (s)*. US Census Bureau, Center for Economic Studies Suitland, MD, USA.

Moses. O Omopekunola, Anna. M. Gorbacheva, and Monday. A. Onoja. In press. Evaluating LLM-generated assessment items: A JBI-guided critical appraisal and checklist. *International Journal of Evaluation and Research in Education*.

OpenAI. 2026. GPT-5.6 terra model documentation. `https://developers.openai.com/api/docs/models/gpt-5.6-terra`. Accessed 2026-08-09.

Qwen Team. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.

Yishen Song, Junlei Du, and Qinhua Zheng. 2025. Automatic item generation for educational assessments: a systematic literature review. *Interactive Learning Environments*, 33(9):5386–5405.

Bin Tan, Nour Armoush, Elisabetta Mazzullo, Okan Bulut, and Mark J. Gierl. 2025. A review of automatic item generation techniques leveraging large language models. *International Journal of Assessment Tools in Education*, 12(2):317–340.

Zhipeng Yin, Sribala Vidyadhari Chinta, Zichong Wang, Matthew Gonzalez, and Wenbin Zhang. 2025. Fairaied: Navigating fairness, bias, and ethics in educational ai applications. *Preprint*, arXiv:2407.18745.

Appendix A Example Source Item and User-Prompt

Part of the user-prompt for generating a multiple-choice item. The source item follows a strict HTML format embedded within JSON.


```
"english": {
    "options": [
      {
        "isCorrectAnswerKey": false,
        "text": "<html><head></head><body><p>91</p></body></html>"
      },
      {
        "isCorrectAnswerKey": false,
        "text": "<p>186</p>"
      },
      {
        "isCorrectAnswerKey": false,
        "text": "<p>1801</p>"
      },
      {
        "isCorrectAnswerKey": true,
        "text": "<p>1860</p>"
      }
    ],
    "prompt": "<p>There are 31 days in May. Calvin does 60 pushups every day of the month of May.</p
        ><p><br/>How many total pushups does he do in the month of May?</p>"
    },
    "scoring": {
    "defaultScore": 0,
    "rubric": "D"
}
```


```
Now, generate 10 new, original items for grade 4 students that follow the same structure and style.
The standards covered in these items is: 'Use the four operations with whole numbers to solve
problems.' Follow the content specifications and fairness guidelines as constraints. Make the new
items similar difficulty as the source item. Otherwise, be creative in the items you generate. The
item content goes in the 'english' JSON field. Instructions on scoring the item goes in the 'scoring'
 JSON field. Return a JSON array structured similarly to the source.
```

Appendix B AI Reviewer Prompt

```
_REVIEW_SYSTEM = (
    "You are a Smarter Balanced assessment content reviewer. "
    "Evaluate generated items strictly and concisely. "
    "Return only valid JSON - no markdown fences, no extra text."
)

_REVIEW_TEMPLATE = """
Review this generated assessment item against the content specification.

Content Specification:
{content_spec}

Fairness Guidelines:
{fairness_guidelines}

Generated Item:
{item_json}

Evaluate whether this item:
1. Matches the target claim, target, and depth of knowledge (DOK)
2. Has language and complexity appropriate for the grade level
3. Has correct item structure (e.g., right number of answer options and correct answers)
4. Avoids fairness or sensitivity issues

Return ONLY this JSON object:
{{"decision": "accept" | "flag" | "reject", "reasoning": "<1-2 sentences>"}}

- "accept": item meets all requirements
- "flag": minor issues that need human review before use
- "reject": significant structural or content problems
"""
```

Using Generative AI to Simulate Item Responses by Skills Insight Score Band

Jianshen Chen[1], Chansoon Lee[2], Kylie Gorney[3], Judit Antal[1]
[1]College Board, [2]American Board of Internal Medicine, [3]Michigan State University

Abstract

This study evaluates GPT-5.4 Thinking simulations of reading-item responses using Skills Insight score-band descriptions and item content. Simulated and empirical item statistics and IRT parameters were compared. Difficulty recovery was strongest, with moderately strong b-parameter correlations, while a and c recovery was weaker, supporting preliminary difficulty evaluation before field testing.

1 Introduction

Field testing has traditionally been used to evaluate the psychometric properties of newly developed items, including difficulty, discrimination, and potential bias. However, field testing is time-consuming, costly, and resource-intensive, motivating interest in alternative approaches.

Text-based methods have been used to predict item difficulty (for a recent review, see Peters et al., 2025) and other item properties. Although promising, these methods produce predicted statistics, such as classical item statistics and item response theory (IRT) parameters, rather than response-level data. More recently, researchers have used large language models (LLMs) to simulate item responses (e.g., Benedetto et al., 2024; Demirkaya et al., 2025; Maeda, 2025; Tang et al., 2026). Because simulated data can be analyzed similarly to field-test data, they may provide richer information before operational use.

2 Purpose of the Study

When simulating item responses, the ability distribution should reflect the target population. This study uses Skills Insight, developed by content specialists, which describes the knowledge, skills, and abilities (KSAs) expected at different score bands. Unlike prior studies that use generic student levels (e.g., Benedetto et al., 2024), IRT ability levels (Demirkaya et al., 2025; Maeda, 2025), or cognitive diagnosis model mastery profiles (Tang et al., 2026), Skills Insight score bands describe what students in specific score ranges are likely to know and do. This allows evaluation of ordered performance across score bands, alignment between item correctness, rationales, and expected skills, and preservation of item- and score-level patterns.

This study investigates whether LLMs can emulate examinees across score bands by generating item responses and brief rationales consistent with Skills Insight KSAs and comparable to operational data. The goal is to evaluate whether Skills Insight–anchored and LLM-generated response data can reproduce empirical patterns in classical item statistics and IRT parameter estimates. If agreement is acceptable, content specialists could use LLMs to obtain early indications of item difficulty and likely responses across score bands.

3 Simulation Study

3.1 Data

Three sources of information are provided to GPT:

1. Skills Insight KSA descriptions for seven score bands on a reading assessment.
2. An empirical ability distribution that specifies the proportion of examinees within each of the seven score bands.
3. Item content, including stems, response options, answer keys, and expert-assigned content categories.

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 140–144
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

The study uses 50 reading items from a fixed form.

3.2 Methods

GPT-5.4 Thinking (OpenAI, 2026) is used to predict whether each simulated student in a specific Skills Insight band is likely to answer each item correctly and generate item scores (i.e., 0 for incorrect or 1 for correct). For each run, the model allocates students across the seven score bands and predicts item responses based on the alignment between cumulative band-level KSAs and the cognitive and linguistic demands implied by item content and expert-labeled categories. The sample prompt and output are shown in Appendices A and B, respectively.

Simulations are conducted under three sample sizes per item: N = 500, 1,000, and 2,000. Each sample size is run with 10 independent replications using the same prompt and input files. Two custom GPTs are created: one without training and one with training. Note that in this study, "training" refers to task-specific refinement using empirical item information. The underlying GPT model parameters are not updated. The 50 items from a fixed form were split into a training set of 30 items and a validation set of 20 items. In the without training scenario, GPT uses the provided files and prompt (see Appendix A) to generate responses and rationales for the 20 items in the validation set directly. In the training scenario, for the 30 training items, empirical IRT parameters are provided, and GPT compares the content-based predicted probability of a correct response with the IRT parameter-based probability. When the absolute discrepancy exceeds a user-set criterion (i.e., 0.20 in this study), GPT refines its content model and prediction. The refined content-based approach is then applied to the 20 validation items, for which the empirical IRT parameters are not provided to GPT. The same prompt shown in Appendix A is used for generating predicted item responses of the 20 validation items.

Classical item statistics, including proportion correct (p) and point-biserial correlation (r_{pbis}), are calculated from GPT-predicted responses, and IRT parameters are estimated based on GPT-predicted responses for the three-parameter logistic (3PL) model using marginal maximum likelihood estimation via the *mirt* package in R (Chalmers, 2012). Results are compared with empirical-data-based classical item statistics and IRT parameter estimates.

3.3 Results

Two conditions the study examines are sample size per item (500, 1000, and 2000) and GPT training status (Without Training and With Training). For each condition, reported values represent means across 10 simulation runs. Table 1 summarizes proportion correct (p), point-biserial correlation (r_{pbis}), and IRT a, b, and c parameter estimates based on the GPT-generated data. The figure shows that the means and variability of p, a, and b from GPT-generated responses generally seem to be reasonable. For example, at N = 1,000 without training, the means (SDs) of p, a, and b are 0.51 (0.18), 1.31 (0.32), and -0.05 (0.87), respectively. In contrast, r_{pbis} and c show relatively small means and variability. For example, under the same condition, the mean (SD) of r_{pbis} and c are 0.45 (0.06) and 0.01 (0.02), respectively. In general, there are minor variations across sample sizes and GPT training status. The only exception to this result is for a. For example, at N = 1,000, the mean (SD) of a increases from 1.31 (0.32) without training to 1.47 (0.48) with training.

		Without Training				With Training			
Statistic	*N*	*Mean*	*SD*	*Min*	*Max*	*Mean*	*SD*	*Min*	*Max*
p	500	0.50	0.19	0.15	0.86	0.51	0.20	0.09	0.93
	1000	0.51	0.18	0.14	0.88	0.53	0.19	0.17	0.91
	2000	0.53	0.18	0.12	0.86	0.51	0.20	0.05	0.92
r_{pbis}	500	0.44	0.05	0.24	0.57	0.47	0.08	0.25	0.63
	1000	0.45	0.06	0.24	0.62	0.46	0.07	0.26	0.63
	2000	0.44	0.06	0.28	0.57	0.44	0.08	0.16	0.62
a	500	1.24	0.24	0.69	2.12	1.45	0.43	0.74	2.75
	1000	1.31	0.32	0.56	2.51	1.47	0.48	0.66	3.02
	2000	1.30	0.28	0.75	2.17	1.34	0.41	0.82	2.51
b	500	0.02	0.91	-2.10	2.24	-0.07	0.94	-2.98	2.14
	1000	-0.05	0.87	-2.64	1.82	-0.15	0.89	-2.57	1.76
	2000	-0.11	0.86	-2.19	1.65	-0.05	1.00	-2.37	3.51
c	500	0.01	0.01	0.00	0.07	0.00	0.01	0.00	0.06
	1000	0.01	0.02	0.00	0.17	0.02	0.03	0.00	0.16
	2000	0.02	0.04	0.00	0.21	0.01	0.01	0.00	0.06

Table 1: Descriptive Statistics

GPT-based item statistics and parameter estimates are evaluated via the Pearson correlation, bias, root mean squared difference (RMSD), and

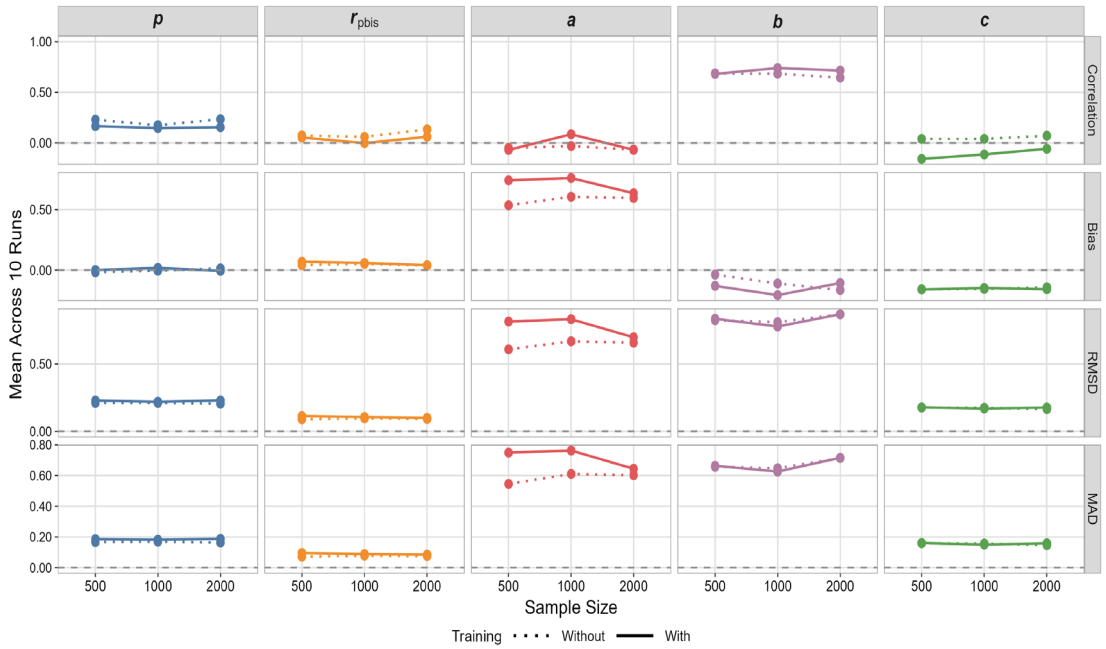


Figure 1: Measures for Classical Item Statistics and IRT Parameter Estimates Without and With Training

mean absolute difference (MAD) using the empirical data-based item statistics and parameter estimates as the benchmark.

Figure 1 presents these measures for the classical item statistics and the IRT item parameter estimates. Increasing sample size and training GPT had minimal impact on performance. Correlations for r_{pbis}, a, and c are very weak, indicating limited agreement between GPT-generated and empirical estimates. In addition, discrimination parameters (r_{pbis} and a) show positive bias, suggesting that GPT-generated data tend to inflate the degree to which items differentiate examinees by ability.

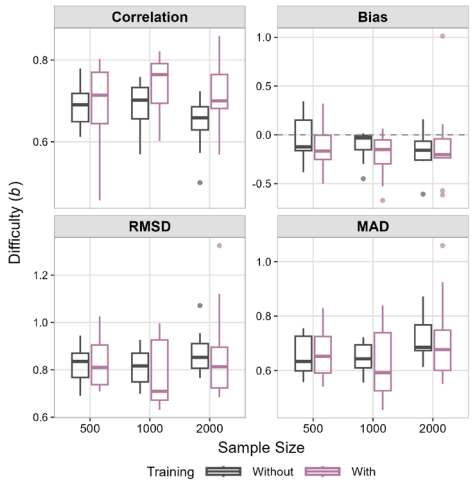


Figure 2: Measures for Item Difficulty Parameter Estimates Without and With Training

In contrast, correlations for item difficulty (p and b) are higher, with means of at least 0.15 and 0.65, respectively. For example, the correlation for b ranges from 0.46 to 0.86, and its association tends to become slightly stronger with training, as shown in Figure 2. Estimates for b show a slight negative bias, indicating that items appear slightly easier in GPT-generated data.

Computation time (in seconds) is illustrated in Figure 3, which tends to increase with sample size. For example, in the without training condition, the mean computation time increases from 302 seconds (SD = 67) at N = 500 to 477 seconds (SD = 156) at N = 2,000. No systematic patterns were observed across the different training conditions.

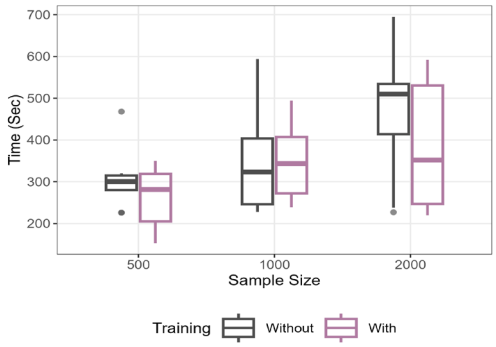


Figure 3: Computation Time Across 10 Runs

4 Discussions and Conclusion

This study provides evidence that Skills Insight–based prompting can be used with generative AI to simulate item responses and approximate some psychometric characteristics of reading items. Across conditions, the strongest agreement was observed for item difficulty, particularly for IRT difficulty (*b*) parameter estimates, suggesting that AI-generated response data can meaningfully differentiate easier from more difficult items when grounded in operationally defined score-band descriptions. In contrast, recovery of discrimination-related estimates (r_{pbis} and *a*) and the *c* parameter was weak. Taken together, the findings suggest that the current approach is more useful for preliminary item-difficulty evaluation than for recovering the full set of psychometric item parameters.

Two findings were unexpected: the with training condition did not consistently outperform the without training condition, and increasing simulated sample size did not materially improve agreement with empirical estimates. These results should be interpreted in light of several limitations. First, as noted earlier, "training" in this study refers to task-specific refinement rather than parameter-level fine-tuning of the underlying LLM. In the with training condition, empirical IRT information from the training items was used to compare content-based predicted probabilities with IRT-based probabilities and to refine the model's prediction process when discrepancies exceeded a prespecified threshold. The model weights themselves were not updated. Thus, the with training condition is better interpreted as an iterative calibration or feedback-based refinement procedure, and the lack of consistent improvement should not be taken as evidence that formal model fine-tuning would be ineffective. Second, the study examined only GPT-5.4 Thinking, so the findings may not generalize to other models or model versions. Third, the simulation was limited to one reading assessment and one fixed form, which constrains generalizability across content areas, test designs, and examinee populations. Future research could replicate the study across models and assessments, and investigate prompting or calibration approaches that improve recovery of the *a* and *c* parameters.

Overall, the findings from this study highlight the potential of combining content-based student skill profiles with LLMs to support early-stage item evaluation. By leveraging existing Skills Insight frameworks, the proposed approach offers a practical and interpretable method for obtaining preliminary indications of item difficulty prior to field testing. Although additional research is needed to strengthen the prediction of item parameters, the overall results demonstrate a promising proof of concept and contribute to the growing body of research on AI-assisted assessment development. As generative AI capabilities continue to advance, approaches such as this may become increasingly valuable as a complement to traditional psychometric methods and item development practices.

References

Benedetto, L., Aradelli, G., Donvito, A., Lucchetti, A., Cappelli, A., & Buttery, P. (2024). Using LLMs to simulate students' responses to exam questions. Findings of the Association for Computational Linguistics: EMNLP 2024, 11351–11368. https://doi.org/10.18653/v1/2024.findings-emnlp.663

Chalmers, R. P. (2012). mirt: A multidimensional item response theory package for the R environment. *Journal of Statistical Software*, *48*(6), 1–29. https://doi.org/10.18637/jss.v048.i06

Demirkaya, O., Wei, H.-R., & Johnson, E. (2025). Simulating rating scale responses with LLMs for early-stage item evaluation. Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 1: Full Papers, 385–392.

Maeda, H. (2025). Field-testing multiple-choice questions with AI examinees: English grammar items. *Educational and Psychological Measurement*, *85*(2), 221–244. https://doi.org/10.1177/00131644241281053

OpenAI. (2026). ChatGPT (GPT-5.4 Thinking) [Large language model]. https://chat.openai.com

Peters, S., Zhang, N., Jiao, H., Li, M., & Zhou, T. (2025). Review of text-based approaches to item difficulty modeling in large-scale assessments. Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 3: Coordinated Session Papers, 37–47.

Tang, X., Lu, Y., Behrens, J. T., & Cheng, Y. (2026). *Can large language models emulate human performance on educational assessments?* PsyArXiv. https://doi.org/10.31234/osf.io/btja8_v2

Appendix A. Prompt for Generating Item Responses and Rationale

Project: Predicted Item Response based on Item Content and Skills Insights only

Use item content, item content constraints in the uploaded file [Constraint File.xlsx], and the knowledge, skills and abilities (KSAs) as described in Skills Insight statements only to predict item responses for n=[500/1000/2000] students. For item matching across files, use item code in the [Item File.json] file to match with item ID in the file [Constraint File.xlsx].

Percentage of students in each score band of the Skills Insights follows the percentages in the Reading excel file uploaded to this Custom GPT. For the simulated students in the same score band, varying the predicted responses considering different extent of mastering KSAs because of different scale scores (e.g., a student with scale score of 239 would be expected to master KSAs slightly better than a student with scale score of 231). No scale score to theta mapping is needed. The prediction is based on whether students with different KSAs are more likely to answer each item correctly or incorrectly based on KSAs described in the Skills Insight statements for each score band and the item content/content constraints. No IRT model is involved.

Include the response matrix, the simulated students' scale scores, the features each item measure, and the rationale of why each simulated student likely answered each item correctly or incorrectly. Produce the output as a downloadable excel file.

Appendix B. Sample Output of Predicted Item Reponses and Rationale

StudentID	ItemCode	IsCorrect	Rationale
S0001	Item XX	0	Likely incorrect because the item depends on using context to determine academic vocabulary, which is above the student's likely demonstrated profile; the student's scale score 206 in band 200–220 indicates beginning foundational reading skills; often struggles unless information is very explicit. Simulated P(correct)=0.020.
S0984	Item XX	1	Likely correct because the item depends on using context to determine academic vocabulary, which is near the edge of the student's likely demonstrated profile; the student's scale score 290 in band 271–300 indicates can draw subtle or complex inferences, handle relatively uncommon academic vocabulary, and make subtle/complex cross-text connections. Simulated P(correct)=0.559.

Exploring AI-driven Methods for Pre-calibrating Difficulty of Mathematics Items with Visual Content

Leah Walker **Jinnie Choi**
Savvas Learning Company
{leah.walker, jinnie.choi}@savvas.com

Abstract

Use of AI is bringing efficiencies in replacing previously laborious field- testing with scaled item parameter pre- calibration, especially in language assessment. However, gaps in research exist in replicating this accomplishment for mathematics assessment items with visual content. This study explores the methodological options of using AI in pre-calibrating multimodal mathematics items.

1 Background

Item pre-calibration, typically done through field-testing, is required for precise, targeted and comparable measurement. Recent advances in artificial intelligence (AI) and natural language processing (NLP) have prompted researchers to investigate whether new methods can automate or replace aspects of costly human field-testing.

Prior work has explored predicting item difficulty directly from linguistic and textual features, including syntactic complexity, readability indices, semantic embeddings, and transformer-based representations (AlKhuzaey et al., 2024; Benedetto et al., 2023). Recent studies proposed an alternative approach in which, rather than directly predicting item parameters, simulated response data were generated using transformer-based large language models (LLMs) and subsequently analyzed using conventional IRT calibration procedures (Feng et al., 2025; Maeda, 2025). This approach preserved important aspects of traditional field-testing workflow while leveraging AI to approximate human response behavior.

Meanwhile, existing AI-based field-testing research has focused primarily on text-dominant reading and language assessment domains. Comparatively little attention has been given to mathematics assessments, where item meaning is often conveyed through symbolic notation, graphs, diagrams, tables, and other visual representations that are integral to the reasoning process itself.

Therefore, before using AI for pre-calibrating these items, alternative methods demand consideration. For example, different strategies for representing visual content as textual input require attention, as transformer-based LLMs are principally designed to process sequential text input. Additionally, it is important to ensure that the representational strategies do not introduce unintended semantic variations, and to evaluate whether the transformer-based model with such input will predict response probabilities and item difficulties comparable to the ground truth empirical results.

2 Aims

This study explores the methodological options for using AI in pre-calibrating multimodal mathematics items, guided by the following questions.

RQ1. What are the different strategies to represent mathematics items with visual content as input for subsequent AI processing?

RQ2. When the visual content is transformed for LLM models, how do we ensure that the representation does not introduce accidental drift?

RQ3. What are the options to generate synthetic examinee response probabilities for these items?

RQ4. How can these results be evaluated for psychometric integrity?

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 145–151
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

3 Related Work

3.1 Methods for Encoding Visual Content in Mathematics Items

Transformer Architectures: Literature suggests three primary transformer architectures for processing mathematical items with graphical representations as input. The first type of approach relies on text-based encoder-only LLMs such as BERT (Bidirectional Encoder Representations from Transformers, Devlin et al., 2019) which improves earlier systems by reading words in human sentence bidirectionally for deep learning, and RoBERTa (Liu et al., 2019) which outperforms BERT with disentangled attention mechanism, dynamic masking and larger pre-training corpora. Because these models cannot process visual information, the visual component is first converted into a textual description via image captioning or optical character recognition (OCR), which then is concatenated with the question stem and passed to the model (Lu et al., 2024).

The second type of approach utilizes vision-language models (VLMs) or large multimodal models (LMMs) such Open AI's GPT4 Vision or Google's Bard/Gemini, which transform images via direct extraction of multi-head self-attention features from patches of the original graphical representations (Li et al., 2023; Lu et al., 2024). This approach minimizes error and redundancy by bypassing the image to text translation layer.

The third type of approach is based on deep fusion encoders which processes item text, visual components, and item-level metadata through separate encoders such as DeBERTaV3 (He et al., 2021) and DINOv3 (Siméoni et al., 2005). The resulting representations are fused using attention pooling and differential weighting of the relative importance of the text versus the image, which then passed as features in the IRT models (Maeda & Lu, 2026).

Hyperparameter Tuning: Because calibrated item datasets are significantly smaller in size than standard NLP corpora that are used for model training, preventing overfitting is the primary focus of hyperparameter tuning. For example, studies explored freezing the lower N layers of BERT, extracting embeddings from the frozen model, and feeding them into external regressors like Ridge or Lasso instead of full end-to-end model fine-tuning (Kapoor et al., 2025; Peters et al., 2025; Chen & Chen, 2026). Also, methods like low-rank adaptation (LoRA) and adapter layers have been utilized to adapt larger models to small psychometric datasets without destroying pre-trained language representations (Scarlatos et al., 2025; Chen et al., 2026; Ormerod & Kim, 2026).

Typical loss functions used in literature include root mean squared error (RMSE) and mean absolute error (MAE) for default continuous item parameters, while evaluation metrics such as accuracy, precision, recall, F1 score, area under the curve and cross-entropy were also used when continuous parameters were discretized into ordinal bins for prediction (e.g., Easy, Medium, and Hard) (Chen & Chen, 2026; Kim, Koo & Lim, 2023; Kim, Kim & Jang, 2023). For search strategies, K-fold cross-validation with typically 5- or 10-fold evaluated strictly at the item or passage level is the standard, preventing inflation of performance metrics due to inadvertent leakage of visual elements or text during the training phase (Badran & Preisach, 2025; Štěpánek et al., 2023).

3.2 Mitigating Drift and Bias in Visual-to-Textual Conversion

Among the encoding strategies, using image captions and 'Alt-text' type of visual-to-textual conversions can cause modality bias and accidental drift in item parameters. To prevent the drift caused by representation, recent literature suggests using vision-language transformers to better approximate the actual human task of synthesizing visual information with reading comprehension, and evaluate multimodal items in their native format (Khan, 2026). For high-stakes assessments, automated transformations require human-in-the-loop cognitive audits on construct equivalence and distractor viability (Burke, 2025). In order to evaluate the impact of LLMs' potential drift and bias, such as tendency to generate systematically easy items, researchers examine the benchmark quality measures against the item characteristic curves (ICC) (Zhou et al., 2026) and how the predicted parameters would alter downstream

ability estimation for examinees (Settles et al., 2020).

3.3 Generating Examinee Responses

Instead of mathematically simulating a dataset with 0 or 1 scores, researchers are increasingly prompting LLMs to act as "proxy examinees" to directly generate synthetic test responses. Srivatsa and colleagues (2025) evaluated this paradigm by treating various LLMs as synthetic examinees on mathematics and reading comprehension assessments. The study applied IRT to the LLM-generated responses to map them onto the same ability scale as real examinee populations, finding that without specific persona prompting, LLMs perform as significantly stronger, high-performing examinees than average human examinees. Other researchers addressed this challenge and introduced a framework which forces LLMs to act as examinees of specific abilities under a ground-truth IRT model using direct preference optimization (DPO) (Scarlatos et al., 2025). Recently, Ormerod and Kim (2026) simulated examinee responses to visual mathematics items using a multimodal LLM based on Qwen 3.5 which was fine-tuned to replicate the systematic error patterns and choice probabilities of examinees conditioned on specific latent ability levels.

3.4 Evaluation of Psychometric Integrity

Statistical Comparisons: To evaluate whether the new approaches are psychometrically sound, studies have compared transformer-based predictions to those from traditional methods such as Explanatory Item Response Models (EIRM), Linear Logistic Test Model (LLTM), and machine learning baselines such as support vector machines, random forests, and penalized regression trained on traditional NLP features (Sharpnack et al., 2024). To prove that interpreting the image actually aids the model, researchers also compared multimodal VLM predictions against text-only and vision-only predictions and examined whether the visual components contribute unique predictive variance to the item parameter beyond what the text provides (Khan, 2026; Maeda & Lu, 2026).

Evaluation of Reconstructed Items: To evaluate how well these models predict parameters or reconstruct new items matching specific targets, the literature typically examines direct parameter recovery using RMSE and MAE that measure the overall magnitude of the prediction error, and rank correlation coefficients such as Spearman's r or Pearson's r to evaluate the relative ordering of the predicted item difficulties (Kolesnikova et al., 2026; Razavi & Powers, 2025).

4 Methods

To evaluate how representation and structural abstraction of mathematics items with visual components influence item parameter prediction and response surface estimation, we developed an empirical workflow combining techniques identified in the related works and predictive modeling, aligned with our four core research questions.

For RQ1, we examine two representation strategies to convert visual mathematical content into text. 'Alt-text' is accessibility-oriented and focused on describing what is visible. 'UI-like' description is what the graph means in the context of the mathematical task. We first convert the visual content into each representation format and then convert them into dense vector embeddings using Sentence-Transformer. These embeddings are then concatenated with item-level metadata to form the input feature set.

For RQ2, to ensure that any differences in psychometric outcomes are caused by the chosen representation strategy rather than unrelated variance introduced by the LLM, we provide textual descriptions to an LLM and prompt it to generate and reconstruct multiple-choice items preserving semantic equivalence with the original assessment items. We verify if the reconstructed items maintained the correct answer key, intended mathematical relationships, and plausible distractors. Items that fail these consistency checks are excluded to prevent them from skewing the analysis.

For RQ3, we explore transformer-based text-only response prediction models using DeBERTa, which estimates response probabilities conditional on structured item representations and simulated examinee ability. We train separate models for each representation condition. This isolation allows the model to

learn the distinct response patterns associated with the level of abstraction and structure unique to each strategy.

For RQ4, psychometric fidelity is evaluated by comparing the transformer-based text-only synthetic response generation and Rasch-generated response probabilities. After reconstructing items, generating response probabilities, and generating responses, we are performing new Rasch calibrations on those responses and comparing the recovered item difficulty parameters directly against the original, pre-calibrated item parameters using RMSE and Pearson correlation to measure how closely they align.

Sample: The study utilizes 330 multiple-choice mathematics assessment items. Each item contains visual information and was previously calibrated using the Rasch model with real examinee data.

To ensure model evaluation reflects the ability of the model to generalize to unexposed items, the calibrated item bank is partitioned into mutually exclusive training and test sets (80%/20%): The training set for optimizing the parameters of the transformer-based model, and the test set for evaluating predictive performance and psychometric approximation accuracy. The split is performed at the item level to ensure that all response patterns associated with a given item are contained within a single partition, preventing information leakage across conditions.

Examinee ability values are sampled consistently across both partitions to maintain comparability of distributions along the latent trait continuum.

5 Results

For RQ1 and RQ2, visual stimuli were converted into structured 'UI-like' representations (Figure 1). This strategy preserves ordinal structure while converting spatial layout into sequential linguistic form suitable for transformer-based processing.

For RQ3 and RQ4, a synthetic dataset was first constructed using a Rasch-based response function over a grid of examinee ability values. Item difficulty parameters were combined with simulated ability to generate probabilistic response values according to the Rasch model.

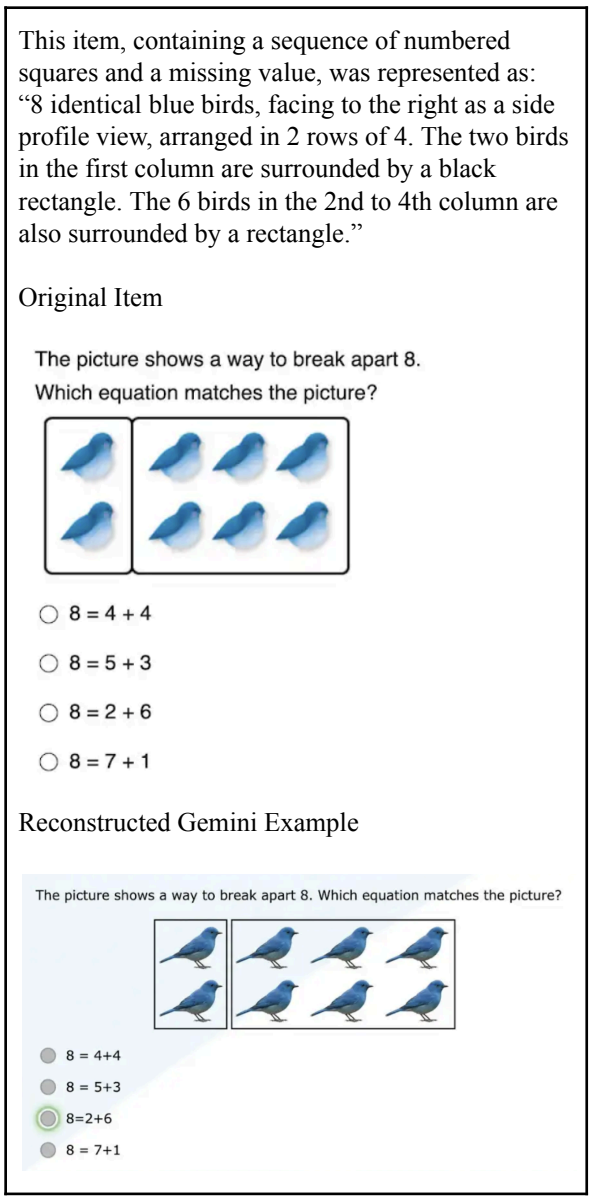


Figure 1. Original Item and Item Reconstructed by LMM (Gemini)

Then, as preliminary baselines, ridge regression models were trained to predict Rasch-generated response probabilities using examinee ability and item embeddings across three prompt structures on a held-out test set (20%). Incorporating explicit UI-like representations achieved the strongest correlation with the response probabilities (RMSE = 0.395, $r = 0.923$).

Because full UI-like representations were not yet available for the entire item set, this baseline reflects a complete-case evaluation of a subset of items. Across the broader text representations, omitting image alt-text yielded better performance (RMSE = 0.199, $r = 0.619$) than including alt-text alongside the stem and options

(RMSE = 0.209, r = 0.579). Overall these preliminary findings reveal a possible tradeoff: standard textual representations provide superior absolute probability prediction, whereas explicit visual and layout descriptions capture correlations more effectively (r = 0.923).

6 Discussion

Results of this ongoing study provide initial evidence regarding the psychometric fidelity of methodological options when pre-calibrating mathematics items with visual content. Specifically, these preliminary findings illustrate how representational abstraction influences item parameter recovery and suggest that transformer-based text-only response models can approximate psychometric response behavior without relying on direct vision-language modeling architectures.

Initial results highlight distinct trade-offs in response probability estimation across the prompt structures. The high Pearson correlation observed under explicit UI-like representations (r = 0.923) demonstrates strong linear alignment between predicted and true response probabilities, even though point-estimation error was higher (RMSE = 0.395). Because full UI-like representations were evaluated on a complete-case subset due to missing layout data, ongoing work will examine whether scaling these representations across the full item bank helps reduce absolute error while preserving this high correlation.

Across the broader text representations, the preliminary finding that omitting alt-text yielded superior probability predictions (RMSE = 0.199, r = 0.619) compared to including alt-text (RMSE = 0.209, r = 0.579) points to an important feature-engineering consideration. This suggests that uncurated alt-text strings may introduce semantic noise into the embedding space, whereas concise textual stems and options offer better calibrated response probability estimates. Ongoing work is exploring how visual metadata can be better structured to improve overall prediction accuracy.

Compared to traditional empirical field-testing methods, which require administering uncalibrated items to large samples of human examinees, this embedding-based pre-calibration framework offers a rapid alternative. Traditional item response theory parameter estimation is slow, costly, and dependent on acquiring adequate sample sizes per item. By contrast, combining item embeddings with examinee ability parameters enables immediate, pre-operational response surface estimation prior to live administration. While synthetic pre-calibration is not intended to replace operational calibration entirely, it provides an efficient screening tool for automated item generation, item bank auditing, and preliminary form assembly.

A critical consideration in AI-assisted psychometric workflows is maintaining item security. Exposing operational assessment content to third-party commercial LLMs risks data leakage and model memorization, where proprietary items could be retained and exposed to examinees in generative contexts. To ensure security, our approach relies on local infrastructure execution and non-invertible feature extraction. Embedding generation is executed entirely on secure private infrastructure using open-source encoder models, preventing data transmission to external APIs or third-party training corpora. Furthermore, items and visual components are converted into fixed-dimension numerical vectors rather than processed through generative text-to-text fine-tuning. These dense vectors capture the latent semantic and structural features necessary for probability prediction while preventing the reconstruction or memorization of raw item text and answer options.

More broadly, by extending AI-based field-testing research into multimodal mathematics assessment contexts, this study contributes to emerging work at the intersection of educational measurement, NLP, and AI-driven assessment development.

References

Samah AlKhuzaey, Floriana Grasso, Terry R. Payne, and Valentina Tamma. 2024. Text-based question difficulty prediction: A systematic review of automatic approaches. *International Journal of Artificial Intelligence in Education*, *34*(3), pages 862–914.

Yahya Badran and Christine Preisach. 2025. Enhancing knowledge tracing through leakage-free and recency-aware embeddings. 2025. *arXiv preprint arXiv:2508.17092*.

Lucca Benedetto, Paolo Cremonesi, Andrea Caines, Paula Buttery, Andrea Cappelli, Andrea Giussani, and Roberto Turrin. 2023. A survey on recent approaches to question difficulty estimation from text. *ACM Computing Surveys*, *55*(9), pages 1–37.

Charles MacDonald Burke. 2025. AI-assisted exam variant generation: A human-in-the-loop framework for automatic item creation. *Education Sciences, 15*(8), page 1029.

Han Chen, Ming Li, Hong Jiao, and Tianyi Zhou. 2026. Representing visual evidence for item difficulty prediction: Visual textualization and image-native modeling. *arXiv preprint arXiv:2608.04554*.

Shi-Ting Chen and Jinsong Chen. 2026. From text to parameters: Predicting item parameters from embedding regularization with reliability and design ceilings. *arXiv preprint arXiv:2607.07141*.

Onur Demirkaya, Hsin-Ro Wei, and Evelyn Johnson. 2025. Simulating rating scale responses with LLMs for early-stage item evaluation. In *Proceedings of the Artificial Intelligence in Measurement and Education Conference: Full Papers*, pages 385–392.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, 1 (Long and short papers)*, pages 4171–4186.

Wanyong Feng, Peter Tran, Stephen Sireci, and Andrew S. Lan. 2025. Reasoning and sampling-augmented MCQ difficulty prediction via LLMs. In *Proceedings of the International Conference on Artificial Intelligence in Education*, pages 31–45. Cham: Springer Nature Switzerland.

Pengcheng He, Jianfeng Gao, and Weizhu Chen. 2021. DeBERTaV3: Improving DeBERTa using electra-style pre-training with gradient-disentangled embedding sharing. *arXiv preprint arXiv:2111.09543*.

Radhika Kapoor, Sang T. Truong, Nick Haber, Maria Araceli Ruiz-Primo, and Benjamin W. Domingue. 2025. Prediction of item difficulty for reading comprehension items by creation of annotated item repository. *arXiv preprint arXiv:2502.20663*.

Samin Khan. 2026. Using vision+language models to predict item difficulty. *arXiv preprint arXiv:2603.04670*.

Gun Il Kim, Sungtae Kim, and Beakcheol Jang. 2023. Classification of mathematical test questions using machine learning on datasets of learning management system questions. *Plos one, 18*(10): e0286989.

Jinsung Kim, Seonmin Koo, and Heuiseok Lim. 2023. A multi-faceted exploration incorporating question difficulty in knowledge tracing for English proficiency assessment. *Electronics, 12*(19), page 4171.

Diana Kolesnikova, Kirill Fedyanin, Abe D. Hofman, Matthieu JS Brinkhuis, and Maria Bolsinova. 2026. Estimating item difficulty with large language models as experts. *arXiv preprint arXiv:2605.18562*.

Yiran Li, Junpeng Wang, Xin Dai, Liang Wang, Chin-Chia Michael Yeh, Yan Zheng, Wei Zhang, and Kwan-Liu Ma. 2023. How does attention work in vision transformers? A visual analytics attempt. *IEEE Transactions on Visualization and Computer Graphics 29*(6), pages 2888–2900.

Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A robustly optimized BERT pretraining approach. arXiv preprint arXiv: 1907.11692.

Pan Lu, Hritik Bansal, Tony Xia, Jiacheng Liu, Chunyuan Li, Hannaneh Hajishirzi, Hao Cheng, Kai-Wei Chang, Michel Galley, and Jianfeng Gao. 2024. MathVista: Evaluating mathematical reasoning of foundation models in visual contexts. In *International Conference on Learning Representations 2024*, pages 23439– 23554.

Hokada Maeda. 2025. Field-testing multiple-choice questions with ai examinees: English grammar items. *Educational and Psychological Measurement*, *85*(2):221–244.

Hokada Maeda and Yikai Lu. 2026. Multimodal test item parameter prediction from text, images, and metadata: Fusing together AI vision and language models. *Educational and Psychological Measurement*, pages 1–25.

Christopher Ormerod and YoungKoung Kim. 2026. Multimodal item parameter estimation using simulated response probabilities. *arXiv preprint arXiv:2608.10154*.

Sydney Peters, Nan Zhang, Hong Jiao, Ming Li, Tianyi Zhou, and Robert Lissitz. 2025. Text-based approaches to item difficulty modeling in large-scale assessments: A systematic review. *arXiv preprint arXiv:2509.23486*.

Pooya Razavi and Sonya Powers. 2006. Estimating item difficulty using large language models and tree-based machine learning algorithms. *International Journal of Artificial Intelligence in Education*: 100015.

Alexander Scarlatos, Nigel Fernandez, Christopher Ormerod, Susan Lottridge, and Andrew Lan. 2025. SMART: Simulated students aligned with item response theory for question difficulty prediction. *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 25082-25105.

Burr Settles, Geoffrey T. LaFlair, and Masato Hagiwara. 2020. Machine learning–driven language assessment. *Transactions of the Association for Computational Linguistics, 8*, pages 247–263.

James Sharpnack, Phoebe Mulcaire, Klinton Bicknell, Geoff LaFlair, and Kevin Yancey. 2024. AutoIRT: Calibrating item response theory models with automated machine learning. *arXiv preprint arXiv:2409.08823*.

Oriane Siméoni, Huy V. Vo, Maximilian Seitzer, Federico Baldassarre, Maxime Oquab, Cijo Jose, Vasil Khalidov, Marc Szafraniec, Seungeun Yi, Michaël Ramamonjisoa, et al. 2025. DINOv3. *arXiv preprint arXiv:2508.10104*.

KV Aditya Srivatsa, Kaushal Kumar Maurya, and Ekaterina Kochmar. 2025. Can LLMs reliably simulate real students' abilities in mathematics and reading comprehension? In *Proceedings of the 20th Workshop on Innovative Use of NLP for Building Educational Applications*, pages 988–1001.

Lubomír Štěpánek, Jana Dlouhá, and Patrícia Martinková. 2023. Item difficulty prediction using item text features: Comparison of predictive performance across machine-learning algorithms. *Mathematics, 11*(19), page 4104.

Hongli Zhou, Hui Huang, Ziqing Zhao, Lvyuan Han, Huicheng Wang, Kehai Chen, Muyun Yang et al. 2026. Lost in benchmarks? rethinking large language model benchmarking with item response theory. In *Proceedings of the AAAI Conference on Artificial Intelligence, 40*(41), pages 35085–35093.

Evaluating Multiple Models for Predicting Item Difficulty in a Principled Assessment Design Framework

Alexandra Perez, M. Christina Schneider, Sangdon Lim, Garron Gianopulos, and Kang Xue

Cambium Assessment, Inc.

Abstract

Item difficulty should, in theory, be predictable from features derived from RPLDs, task characteristics, and linguistic complexity. This study evaluates the performance of multiple statistical and machine learning models, linear regression, LASSO regression, random forest, and gradient boosting, in estimating item difficulty. We found similar results across all four models, the item features selected explain 53%–56% of the variance across all grades. Across all four models, number of words in an item, Amount of Text/Source of Inference and grade were the top three important features in predicting item difficulty.

1 Background

Recent advances in assessment design emphasize the integration of cognitive theory, task models, and the use of construct relevant features intended to drive item difficulty (difficulty drivers) to engineer items in particular performance levels (Huff et al., 2025). Within this paradigm, principled assessment design (PAD) links both Range performance level descriptors (RPLDs; Egan et al., 2012) and difficulty drivers to task characteristics to support claims about what students know and can do.

RPLDs are intended to provide descriptions of increasing knowledge, skills, and abilities that are expected of students based on a state's content standards as student abilities increase along a test scale. They are intended to provide content-based interpretations regarding what scale scores within a particular performance level represent within a grade. They are also intended to describe text-task interactions (Schneider et al., 2026) when developed for reading. Text-task interactions are a learning-science-based theory of reading comprehension (Valencia et al., 2017). RPLDs may have different characteristics based on test developers' theory of reading development; however, following the Egan et al. (2012) development framework, RPLDs for English language arts (ELA) optimally describe text-task interactions rather than text complexity. When validating assessments developed under a PAD framework, item difficulty modeling (IDM) is a critical component.

Item difficulty should, in theory, be predictable from features derived from RPLDs, task characteristics, and linguistic complexity. In the future, the goal is for these features to be transparent and useful in supporting item writers and their Large Language Model (LLM) assistants in constructing items with differing characteristics that align to specific areas of the test scale that represent changing stages of student cognition. Range PLDs have been growing in use as a component of item design. Researchers have increasingly applied machine learning techniques and LLMs to predict item difficulty.

Sano (2016) used a linear regression of stem, answer choice options, and passage textual complexity features of 60 Grade 8 NAEP items and found 9 of the 11 predictors that were significant related to NLP features parsed separately for stems, the key, and the options. He highlighted correlations of features that reached

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 152–158
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

.30. Štěpánek et al. (2023) found LASSO regression, ridge regression, and elastic net to perform the best, having the lowest root mean squared error (RMSE) when comparing models. Razavi and Powers (2025) used LLMs to derive item features to predict item difficulty. They conducted an IDM study and found that the random forest model and gradient boosting model produced similar results.

Machine learning techniques have included linear regression (Hastie et al., 2009), random forests (Breiman, 2001), LASSO regression (Tibshirani, 1996), and gradient boosting (XGBoost; Chen & Guestrin, 2016). However, we believe there remains a need to evaluate how different modeling approaches perform within a PAD/IDM framework which values feature transparency to determine which features most strongly predict item difficulty across grade levels in construct relevant ways.

This study investigates the extent to which item characteristics derived from a combination of human coded features and feature extractions predict item difficulty. ELA items have historically been difficult to predict (Ferrara et al., 2022) with R^2 above .50 in a student population. Our aims follow:

1. Evaluate the predictive performance of multiple statistical and machine learning models, including linear regression, random forest, LASSO regression, and gradient boosting (XGBoost), in estimating item difficulty.

2. Identify the most influential item features contributing to difficulty across grades 3–8 in an ELA state assessment.

3. Inform item development practices, with the goal of guiding item writers toward features that predict intended difficulty levels.

2 Data

The data for this study came from items and human coding of those items measuring ELA standards for grades 3–8. Psychometricians used the 2-parameter model for dichotomous items and the generalized partial credit model (GPCM; Muraki, 1992) for polytomous items for vertical scaling. Items were located at the point of the scale where the response probability was .50 (RP50) for the score point or maximum score point for polytomous items. There were 1,264 items, and the number of items per grade ranged from 128 to 288, with grades 4 and 7 having the fewest items.

3 Methods

The task features came from the item banking software, human derived features, the Python Natural Language Toolkit (NLTK; Bird et al., 2009), and Python's textstat library (Shivam, 2024).

To determine the set of predictors included in our model, we used principal components analysis (PCA). However, because many variables make up these components, interpreting them and using the results in practice can be challenging. Therefore, we extracted the components that contributed to 95% of the variance and selected the variable that had the highest contribution in each component. This allowed us to be more confident that the variable importance measures were not influenced by correlated predictors. The final set of predictors included item RPLD level, Amount of Text/Source of Inference (Valencia et al., 2017), Grade, binary item types (multiple select, other, and EBSR), and item text features (number of words, Linsear Write index, Dale Chall index, Coleman Liau index, and whether the response options included a 'paragraph prompt').

We employed separate model-training and model-testing processes for the models, ensuring a more accurate evaluation using a 5-fold cross validation. We split the data into five folds (parts). Each model was trained on four folds and tested on the remaining fold. Because this was repeated five times, each fold was used as the test set once. For model tuning, we repeated this 5-fold cross validation 10 times. For model analysis, we repeated this 5-fold cross validation 100 times.

We compare Ordinary Least Squares (OLS) linear regression, LASSO regression, random forest regression model, and gradient boosting (XGBoost) model. To compare model performance, we examined RMSE and R^2.

For LASSO regression, we initially tuned the model to find the best lambda, as lambda approaches zero, the closer the model is to linear regression (Tibshirani, 1996). For the random forest model, we initially tuned the model to find the best combination of the following hyperparameters:

- number of trees,
- maximum number of splits from the root to the leaf (max depth),
- number of features randomly considered at each split (mtry),
- and the smallest number of observations allowed in a leaf (minimum node size).

Finally, for the gradient boosting model, we initially tuned the model to find the best combination of the following hyperparameters:

- how much each new tree is allowed to adjust the model's predictions (eta),
- how complex each tree can be (max depth),
- how much data is required before a split is allowed (minimum child weight),
- the fraction of rows used for each tree (subsample),
- and the fraction of predictors used for each tree (column sample by tree).

4 Results

Table 1 and Table 2 show the results from the 5-fold cross validation of the linear regression model, the LASSO regression model, the random forest regression model, and the gradient boosting model. For the individual grade-level analyses, we used the predictions from the testing data and calculated R^2 separately for each grade by subsetting the results by grade level. Linear regression and LASSO regression expect linear relationships whereas random forests and gradient boosting show the influence of interactions of features.

The amount of variance explained and RMSE in the testing data for all four models were similar across all grades and within grades. The item features selected explain 53%–56% of the variance across all grades. The RMSE is approximately 0.8. Additionally, for the R^2 and RMSE, the difference between the testing and training dataset was not large, therefore there is minimal overfitting. The random forest model revealed the greatest degree of overfitting, as evidenced by the largest discrepancy between training and testing performance. The lambda value for the LASSO regression model was 0.004, indicating that this model performed similarly to the linear regression model.

Table 1. Average Variance Explained (R^2) Comparative Results

	Average R^2 of Training and Testing Data							
	Training Data	Testing Data						
	All Grades	All Grades	Grade 3	Grade 4	Grade 5	Grade 6	Grade 7	Grade 8
Linear Regression	0.55	0.53	0.53	0.09	0.58	0.38	0.33	0.37
LASSO Regression	0.55	0.53	0.53	0.10	0.58	0.38	0.33	0.37
Random Forest	0.71	0.55	0.59	0.12	0.61	0.36	0.27	0.40
Gradient Boosting	0.67	0.56	0.59	0.11	0.60	0.37	0.31	0.41

Note. The outcome was RP50 of items. We employed a 5-fold CV process and repeated it 100 times.

Table 2. Average Root Mean Squared Error (RMSE) Comparative Results

	Average RMSE of Training and Testing Data							
	Training Data	Testing Data						
	All Grades	All Grades	Grade 3	Grade 4	Grade 5	Grade 6	Grade 7	Grade 8
Linear Regression	0.84	0.85	0.95	0.83	0.82	0.74	0.89	0.82
LASSO Regression	0.84	0.85	0.95	0.83	0.82	0.74	0.89	0.82
Random Forest	0.67	0.83	0.88	0.82	0.78	0.76	0.93	0.80
Gradient Boosting	0.71	0.83	0.87	0.82	0.79	0.75	0.91	0.80

Note. The outcome was RP50 of items. We employed a 5-fold CV process and repeated it 100 times.

The results were not consistent across grades. The variance explained was strongest for Grades 3 and 5 and slightly lower for Grades 6 through 8. Grade 4 was an outlier in comparison to other grades. Grade 4 could be lower because the construct differs, because of the smaller sample size, or a combination. Grades 4 and 7 had lower percentages of items aligned to RPLD Level 4 and Text Source Level 4 compared to other grades.

Table 3 presents evaluation results from the linear regression model. These results represent averages across the 100 5-fold cross-validation datasets across grades. The significant features are highlighted in both models. The results are similar across the two models. All variables are statistically significant except for the item containing a paragraph prompt, Dale Chall index, and Coleman Liau Index. The significant predictors indicate that as the predictor increases, the item difficulty increases, holding all else constant. If an item is an EBSR item, multiple select item, or other item type, the item difficulty increases indicating that multiple choice items predict lower item difficulty for these items. For LASSO regression, similar features were found important in predicting item difficulty, number of words in an item, Amount of Text/Source of Inference, Grade, and the three item type features.

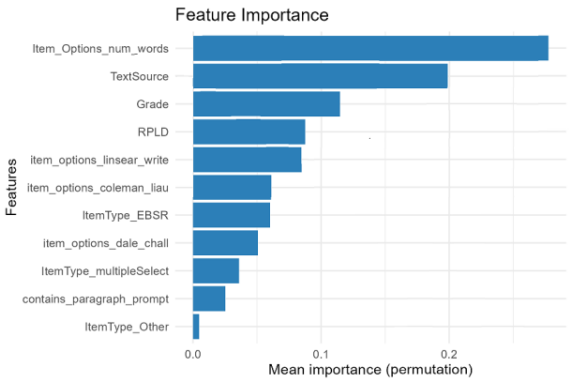
Table 3. Average Regression Coefficients over Replications Across Grades

	b	*p*	CI low	CI high
(Intercept)	-3.21	0.00	-3.59	-2.83
# of Words	0.01	0.00	0.01	0.01
Linsear Write Index	0.03	0.00	0.01	0.04
Contains Paragraph Prompt	-0.04	0.53	-0.17	0.09
Item Type – MS	0.82	0.00	0.63	1.01
Item Type -Other	0.30	0.01	0.09	0.50
Amount of Text Source	0.32	0.00	0.25	0.38
Item Type – EBSR	0.44	0.00	0.25	0.64
Dale Chall Index	0.05	0.08	0.00	0.10
RPLD	0.07	0.04	0.01	0.14
Coleman Liau Index	0.00	0.72	-0.03	0.03
Grade	0.16	0.00	0.12	0.19

Note. The outcome was RP50 of items. We employed a 5-fold CV process and repeated it 100 times. We averaged the coefficients across the training datasets.

Figure 1 shows the feature importance of the random forest model. As shown, number of words, Amount of Text/Source of Inference, Grade, and RPLD are the most important features predicting item difficulty in order of performance. For the gradient boosting model, number of words, Amount of Text/Source of Inference and Grade were the strongest features.

Figure 1: Feature Importance from Random Forest Model



Across all four models, the most important features predicting item difficulty were number of words, Amount of Text/Source of Inference and Grade. For linear regression and LASSO regression, we standardized the coefficients to ensure fair comparison to random forest and gradient boosting. While RPLD is a significant feature in the models, it does not predict an increase in item difficulty as great as Amount of Text/Source of Inference. See Table 4 for the list of important features by model. These features may be contextualized into features related to stem and options reading load and vocabulary, source of evidence in the text needed to respond correctly, RPLDs, and item types.

Table 4. Feature Importance Ordered by Model

Importance	Linear Regression	LASSO Regression	Random Forest	XGBoost
1	Number of Words	Number of Words	Number of Words	Number of Words
2	Amount of Text/Source of Inference	Amount of Text/Source of Inference	Amount of Text/Source of Inference	Amount of Text/Source of Inference
3	Grade	Grade	Grade	Grade
4	Multiple Select Item Type	Multiple Select Item Type	RPLD	Dale Chall Index
5	EBSR Item Type	EBSR Item Type	Linsear Write Index	Multiple Select Item Type
6	Linsear Write Index	Linsear Write Index	Coleman Liau Index	Coleman Liau Index
7	Other Item Type	RPLD	EBSR Item Type	Linsear Write Index
8	Dale Chall Index	Dale Chall Index	Dale Chall Index	RPLD
9	RPLD	Other Item Type	Multiple Select Item Type	EBSR Item Type
10	Coleman Liau Index	Contains Paragraph Prompt	Contains Paragraph Prompt	Other Item Type
11	Contains Paragraph Prompt	Coleman Liau Index	Other Item Type	Contains Paragraph Prompt

5 Conclusion and Discussion

Overall, the variables we included in the model accounted for more than 50% of the variance across all grades. There was minimal overfitting in the models as the training dataset's average R^2 and RMSE were similar for the testing dataset, specifically for the linear regression and LASSO regression. As expected, random forest and gradient boosting had more overfitting, but gradient boosting outperformed random forest in this aspect. Within grades, the average variance explained dropped in grades 4, 6, 7, and 8. For these grades, we could have seen a drop in performance if the construct differs across grades; that is, if the feature combinations were being deployed differently across grades.

Across all four models, number of words in an item, Amount of Text/Source of Inference and grade were the top three important features in predicting item difficulty. For linear regression and LASSO regression, the item type features also contributed to the variation in item difficulty. For the random forest model, the next most important feature was RPLD.

For this study, we found that the four models performed similarly to each other, but the features explaining the variation in item difficulty could differ. If item writers need to craft items that are likely to be more difficult, they should find opportunities to ask questions using text characteristics at the higher end of the Amount of Text/Source of Inference rubric. These types of questions also will tend to involve longer answer choice options and target the higher end of the RPLDs. However, based on the inconsistent finding of RPLDs across models, we recommend that Amount of Text/Source of Inference be a difficulty driver embedded in the RPLDs in the future.

6 References

Bird, S., Klein, E., & Loper, E. (2009). *Natural language processing with Python*. O'Reilly Media.

Breiman, L. (2001). Random forests. *Machine Learning, 45*(1), 5-32. https://doi.org/10.1023/A:1010933404324

Chen, T., & Guestrin, C. (2016). XGBoost: A Scalable Tree Boosting System. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge*

Discovery and Data Mining (pp. 785–794). https://doi.org/10.1145/2939672.2939785

Egan, K. L., Schneider, M. C., & Ferrara, S. (2012). Performance-level descriptors. In G. J. Cizek (Ed.), *Setting performance standards: Foundations, methods, and innovations* (2nd ed., pp. 79–106). Routledge.

Ferrara, S., Steedle, J., & Franz, R. (2022). Response demands of reading comprehension test items: A review of item difficulty modeling studies. *Applied Measurement in Education,35*(3), 237–253.

Hastie, T., Tibshirani, R., & Friedman, J. (2009). The Elements of Statistical Learning: Data Mining, Inference, and Prediction (2nd ed.). Springer.

Huff, K., Nichols, P., & Schneider, M. C. (2025). Designing and developing educational assessments. In L. L. Cook & M. J. Pitoniak (Eds.), *Educational measurement (5th ed*., pp. 441–511). National Council on Measurement in Education. DOI: 10.1093/oso/9780197654965.003.0007

Muraki, E. (1992). A generalized partial credit model: Application of an EM algorithm. *Applied Psychological Measurement, 16*(2), 159–176. https://doi.org/10.1177/014662169201600206

Razavi, P., & Powers, S. J. (2025). Estimating Item Difficulty Using Large Language Models and Tree-Based Machine Learning Algorithms. *10.48550/arXiv.2504.08804.*

Sano, M. (2016 April). Improvements in automated capturing of psycho-linguistic features in reading assessment text. Paper presented at the annual meeting of the National Council on Measurement in Education, Washington, DC.

Schneider, M.C., Chen, J., & Heneger, J. (2026). Text complexity versus task complexity: Item difficulty modeling for reading items. Practical Assessment, Research, & Evaluation, 31(1)(5). https://doi.org/10.7275/pare.2928

Shivam, S. (2024). textstat: A Python package for text statistics (Version 0.7.4) [Computer software]. https://pypi.org/project/textstat/.

Štěpánek, L., Dlouhá, J., & Martinková, P. (2023). Item difficulty prediction using item text features: Comparison of predictive performance across machine-learning algorithms. *Mathematics, 11*(19), 4104. https://doi.org/10.3390/math11194104

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological), 58(*1), 267-288.

Valencia, S. W., Wixson, K. K., Ackerman, T., & Sanders, E. (2017). Identifying Text-Task-Reader Interactions Related to Item and Block Difficulty in the National Assessment for Educational Progress Reading Assessment.

Transformer-Aided Detection of Gaming in Constructed-Response English Language Assessments

Melanie Sharif[1], **Scott Hellman**[2*], **Martha Bellows**[1], **Susan Lottridge**[1]
[1]Pearson plc [2]Macmillan Learning

Abstract

Gaming responses in constructed-response English language assessments can threaten score validity, making robust detection an important function for assessment organizations. We compare handcrafted-feature baselines with frozen transformer embeddings and embedding–feature ensembles for gaming detection across three constructed-response item types under familiar-prompt and unseen-prompt evaluation regimes, and we test whether a single pooled model can replace item-type specialists. Transformer representations improve agreement with gaming labels in most conditions, and an embedding–feature ensemble is the best configuration in every item-type/regime split, or statistically indistinguishable from it, with the best configuration reaching ROC-AUC = 0.95 (Cohen's $\kappa = 0.75$). In the lowest-prevalence item type under unseen prompts, however, the ensemble provides no significant gain over handcrafted features, and frozen embeddings alone are markedly weaker there. A pooled model performs comparably to item-type specialists in the familiar-prompt regime. We also show that these results are bounded by the label-generation process: review process labels record gaming that was detected, while missed gaming is recorded as negative. Reported metrics therefore estimate agreement with the existing review process rather than construct-level detection accuracy. We argue for supplementing standard agreement metrics with reference-conditioned recall and targeted blind re-review of candidate misses.

Keywords: automated scoring, gaming detection, constructed response, transformer embeddings, label quality, validity, label noise

[*] Work done while at Pearson plc

1 Introduction

1.1 Gaming as a validity threat

Automated scoring in constructed-response language assessment relies on the assumption that produced responses provide evidence of the target construct (Mislevy et al., 2003; Williamson et al., 2012). When test-takers substitute memorized templates, copied prompt or passage material, filler, and/or generic text for a construct-relevant response, fooling content models, this is called gaming by practitioners (Hughes and Clesham, 2024). It is also collectively known as construct-irrelevant response strategies, or CIRS (Bejar et al., 2014).

Gaming is thus a direct attack on test validity, as automated scoring systems can reward surface features that resemble valid responses even when the underlying construct evidence is weak or absent (Higgins and Heilman, 2014; Ding et al., 2020). Gaming detection therefore becomes a validity and quality-assurance function. Some detectors rely on handcrafted indicators such as template overlap and prompt or passage copying; these features encode known failure modes but can miss strategies not anticipated in advance (Samant et al., 2025).

Transformer encoders offer a complementary approach by representing semantic and syntactic structure without requiring each gaming pattern to be manually specified, motivating the comparison in this study.

1.2 Label noise and interpretive constraints

Gaming detection differs from many automated scoring tasks because its value lies in identifying gaming that the current review process misses, not merely reproducing existing labels. This creates a measurement problem: a response receives a positive gaming label only if it has already been caught, while gaming missed by review is recorded as negative. The observed labels therefore contain one-sided noise concentrated in the negative class.

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 159–167
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

This is a positive–unlabeled setting: positives are observed only when detection succeeds, and the negative class mixes true negatives with undetected positives (Elkan and Noto, 2008; Bekker and Davis, 2020). A model that flags previously missed gaming may be scored as producing a false positive, even when the flag is useful.

Importantly, conventional accuracy, agreement, precision, and recall consequently estimate agreement with the existing review process rather than pure detection accuracy against the construct. We therefore interpret all performance figures as label-agreement estimates and supplement them with metrics aimed at the deployment question: whether a model recovers positives the reference detector missed. Coverage of missed gaming is the quantity of interest, while κ and AUC serve to rank model quality.

Because test-takers adapt to current detection methods, the target category is adversarial and non-stationary. A detector model's ability to catch latent gaming features is therefore crucial for providing support to human reviewers. Highly sensitive models can help scorers better identify gaming over time, which can gradually improve model training sets and close the coverage gap.

1.3 A note on scoring policies

The corpus does not support an inter-rater reliability (IRR) or "human ceiling" comparison: no response was scored blind by two humans. Instead, the second read for responses in this corpus is a read-behind of the first, so IRR is un-estimated.

The reference labels are also not independent of the systems under evaluation. Adjudication is triggered by a disagreement of more than two score points on a content trait between any two reads, where a positive automated gaming label functions as a content-trait read of score 0. Since the reference detector participates in producing the labels against which new models are scored, agreement with those labels partly reflects shared machinery rather than independent construct evidence (§1.2).

1.4 Paper Contributions

The main contributions of this paper are as follows. First, we evaluate transformer-aided modeling, embedding–feature ensembles, and the handcrafted-feature approach on the gaming detection task across three constructed-response item types and two generalization regimes: familiar prompts and unseen prompts. Second, we evaluate pooled models across item types alongside specialist models to explore whether a pooled model is a defensible choice.

Third, we present an evaluation protocol for labels that are detection outcomes rather than ground truth. This includes reference-conditioned recall that distinguishes between gaming the reference detector caught versus gaming it missed, and prioritizing prevalence-robust agreement metrics (accuracy, Cohen's κ, recall, and ROC-AUC), with ROC-AUC as the primary criterion for both model selection and reporting; precision and F1 are class-dependent, are reported on the balanced condition only, and are not used for selection or cross-condition comparison (§4.1, §4.4).

Finally, we explain why gaming detection is difficult to evaluate under existing labeling conditions and propose concrete, inexpensive remedies that can move this area toward better measurement practice.

2 Related work

2.1 Gaming detection in automated scoring

Automated essay scoring (AES) systems infer a latent writing construct from surface features of a response, and a substantial literature has shown that this inference can be subverted by construct-irrelevant response strategies (CIRS) collectively termed gaming (Bejar et al., 2013, 2014; Higgins and Heilman, 2014; Ding et al., 2020; Hughes and Clesham, 2024; Samant et al., 2025). One of the most conspicuous examples remains the BABEL generator, whose grammatical gibberish nonetheless earned near-perfect scores from production scoring engines (Perelman et al., 2014). This work prompted dedicated detection advisories in response (Cahill et al., 2018), which Perelman (2020) argued did not fully resolve the validity problem. Crucially, the shift from feature-based to deep learning models did not resolve this fragility: neural scorers have been characterized as both overstable and oversensitive, assigning high scores to adversarial inputs produced by repeating, shuffling, deleting, or injecting content and by inserting off-topic or nonsensical text (Lottridge et al., 2020; Ding et al., 2020; Kabra et al., 2020; Singla et al., 2023). As a result of these persistent vulnerabilities, gaming detection remains a vital asset for assessment organizations.

Recent work has focused on detecting templated text as a proxy for inauthentic responses – an ad-

Item type	N	Human-labelled
Type A	111,362	107,789
Type B	637,162	586,661
Type C	167,000	166,564
All	915,524	861,014

Table 1: Corpus size and number of human labels by item type.

versarial strategy distinct from previous ones documented in the literature (Samant et al., 2025). Despite this targeted approach, a gap remains between what handcrafted features like templates and text patterns can catch and the scope of adversarial strategies in use. The sheer volume of linguistic features present in large transformer models, combined with the strength of handcrafted features, could offer a substantial next step in detection.

3 Data

3.1 Assessment context and corpus

Data are drawn from an adult English language proficiency assessment with three constructed-response item types (Type A, Type B, and Type C).

In addition to content scores, every response across item types is subject to an automated gaming review that produces a ternary label — High-Confidence Gaming (GH), Low-Confidence Gaming (GL), and No Gaming (NG). The ternary labels have been collapsed to binary labels for the sake of this study.

3.2 Corpus size

We use a corpus of 915,524 responses, of which 861,014 (94.0%) carry a positive or negative human gaming label. The remaining 6% of responses without a human label are removed from the response pool. For prompt counts and train/test prompt splits, see §4.1.

4 Method

4.1 Splits and generalization regimes

We construct class balanced 50/50 train/test splits per item type, plus a pooled model trained across all three, under two regimes:

- **within_prompt** — random assignment; prompts appear in both train and test. This corresponds to deployment against a known item bank.
- **cross_prompt** — prompt-disjoint, constructed within each item type; no prompt appearing in training appears in test. This corresponds to the cold-start case, a newly authored item before any of its responses have been reviewed.

Regime	Type	train	test	Train gaming	Test gaming
within_prompt	Type A	6,000	2,000	3,000	1,000
within_prompt	Type B	1,756	390	878	195
within_prompt	Type C	6,000	2,000	3,000	1,000
within_prompt	all	5,756	1,722	2,878	861
cross_prompt	Type A	6,000	2,000	3,000	1,000
cross_prompt	Type B	1,620	394	810	197
cross_prompt	Type C	6,000	2,000	3,000	1,000
cross_prompt	all	5,558	1,746	2,779	873

Table 2: Train/test sizes and class label balance by item type and prompt regime.

Splits were requested at 6,000 train / 2,000 test. Type A and Type C achieve this; Type B and the pooled splits fall short because human-labelled positives run out, due to the extremely low prevalence of gaming in that item type (Table 2).

To ensure that enough human-labelled responses are available for inclusion in the test set, the test set is capped at the proportional share the train/test ratio implies, grouped by item type so that a pooled split cannot cover a shortfall in one type by overdrawing another.

Test sets are class balanced 50/50 by construction, standard practice for evaluating classifiers under severe class imbalance (He and Garcia, 2009). We apply the same 50/50 balancing to all three item types so that agreement metrics are computed under a common class distribution across item types and for the pooled model, keeping cross-type and pooled-versus-specialist comparisons on a common footing (Table A1). Prompt-disjoint (cross_prompt) prompt counts are Type A 14/6, Type B 67/14, Type C 20/7, and pooled 95/33 (train/test). Recall and ROC-AUC are class-invariant and transfer to deployment prevalence unchanged.

4.2 Models

There are two baseline model configurations per item type split. The baseline feature set comprises handcrafted features: for Types A and B, type-token ratio, comma density, minimum template distance, and authentic token percentage; for Type C, the number and percentage of authentic tokens over-

all, within templated text, and within prompt text, together with minimum template distance. The extended baseline set adds five runtime classical text features: passage similarity, bigram overlap, entity overlap, key-term coverage, and content-word ratio. Both feature sets are then used in an ElasticNet-regularized logistic regression model.

For the transformer-aided modeling, there are four encoder configurations: BGE-large-en-v1.5 (bi-encoder), longformer-base-4096, longformer-large-4096, and longformer-base-4096 used as a cross-encoder, taking the [CLS] state over a joint prompt–response input. The cross-encoder is included as a deliberate design contrast: the bi-encoders never see prompt and response together, so any gaming signal that consists of a mismatch between the two must be recovered downstream rather than attended to directly. Embeddings are reduced by PCA over {5, 10, 20, 50, 100, 200, 500}, skipping dimensions above the training matrix rank. Each encoder's embeddings are then passed to three classifiers evaluated separately — ElasticNet logistic regression, random forest, and RBF SVM — so that each encoder–classifier pair is a distinct pipeline rather than a combined model.

An ensemble method is also tested by selecting the best embedding source, chosen by nested cross-validation, and concatenating with the handcrafted feature block. PCA is then applied to the embedding block only, and logistic regression is used as the classifier.

4.3 Selection protocol

Nested cross-validation is used for model selection: an inner loop searches hyperparameters within each outer training fold, and the outer loop estimates how the whole selection procedure generalizes (Cawley and Talbot, 2010). All preprocessing — PCA and scaling — is inside the pipeline and refit per fold. The test set is scored exactly once and is never consulted for selection. Model configurations are selected by mean inner-fold ROC-AUC, a threshold-free criterion aligned with the primary reporting metric (§4.4); F1 is not used for selection. The nested cross-validation uses five folds, and random seed 42 is fixed throughout.

4.4 Metrics

Since our labels are review process outcomes rather than an independent measure of the construct (§1.2, §1.3), conventional agreement metrics are necessary but not sufficient. Table 3 summarizes the metrics reported throughout the study and the purpose each serves.

All agreement metrics are reported at each classifier's default decision rule (0.5). No threshold tuning is applied, and no operating-point deployment claims are made. ROC-AUC is the primary metric reported throughout: it is threshold-free — scoring ranking quality across all possible decision thresholds rather than at a single cut — prevalence-robust, and is also the criterion used for model selection (§4.3), so selection and reporting use the same metric.

We additionally evaluate performance using Cohen's κ, which scores each classifier at its default 0.5 decision rule rather than across thresholds. In test sets balanced 50/50 by construction (§4.1), chance agreement is fixed at 0.5 and κ reduces to a linear rescaling of accuracy ($\kappa = 2 \cdot \text{accuracy} - 1$). We report κ rather than raw accuracy because it is the standard in this literature and keeps comparisons legible across item types. Pooled item type slices are also class balanced (Table 7).

5 Results

5.1 Best configuration per split

Table 4 reports the best configuration for each item type and prompt regime, using held-out test data; the longformer-large encoder stood out, appearing in every winning configuration except Type A under familiar prompts. No cross-encoder configuration was selected in any split. The final column shows the change in kappa when moving from familiar prompts to unseen prompts. Table 5 then compares model families using ROC-AUC, the same threshold-free metric used for model selection.

5.2 Test set performance by model family

Across all eight splits, the embedding–feature ensemble improves ROC-AUC over the handcrafted baseline, with gains ranging from $+0.05$ (Type B, cross-prompt) to $+0.22$ (Type C, within-prompt). The ensemble is the best-performing configuration in every split except within-prompt Type B, where it is statistically indistinguishable from the best frozen embedding (Table 5).

5.3 Reference-conditioned recall

Recall on positives the reference detector already flagged is uniformly high (0.94–0.99), whereas recall on positives the reference detector missed is

Metric group	What is reported	Interpretive role
Conventional agreement	Cohen's kappa, F1, precision, recall, and ROC-AUC with 95% percentile bootstrap confidence intervals from 2,000 resamples; McNemar's test for paired model comparisons on identical test items.	Supports model-to-model comparisons within the observed label set. Label bias affects both arms of a paired comparison, so differences remain interpretable even when the labels are imperfect.
Reference-conditioned recall	Recall on positives the reference detector already flagged; recall on positives the reference detector did not flag.	Separates easy agreement with the reference detector from the more valuable case: recovering human-confirmed gaming the reference detector missed.
Unconfirmed flags	Model-positive / label-negative records. Traditional false positives.	Not treated as ordinary false positives, because this cell contains both model errors and possible genuine catches missed by the review process. This is the target population for the blind re-review proposed in §6.

Table 3: Study evaluation metrics by group along with reporting and interpretative information.

lower and more variable (0.57–0.84). The gap is modest under familiar prompts (roughly 0.13–0.17) but widens under unseen prompts for Type A and Type B (to 0.38 and 0.34); even so, the model recovers a nontrivial share of missed gaming in every condition (Table 6).

5.4 Pooled model vs. specialists

The pooled model matches or slightly trails the item-type specialists in every slice, with Cohen's κ differences ranging from 0.000 (Type B, within-prompt) to -0.076 (Type A, within-prompt); the specialist estimate exceeds the pooled slice's bootstrap interval only for within-prompt Type A (Table 7).

The full metric set behind this comparison is reported in Table A1 (Appendix A).

6 Discussion

The first finding is that transformer representations improve agreement with gaming labels in most item types and regimes, but not uniformly. For Type B under cross-prompt evaluation — the lowest-prevalence, most data-constrained condition — the ensemble is statistically indistinguishable from handcrafted features (McNemar p = 0.65). Frozen embeddings therefore provide no significant gain there. Frozen embeddings alone are also markedly weak in this condition (ROC-AUC = 0.7914; $\kappa \approx 0.21$), with the handcrafted block carrying the ensemble (Table 5, Table 6).

The second finding is that unseen prompts are seemingly influenced by the default cutoff rather than ranking quality. ROC-AUC drops are modest and uniform across item types (ΔROC-AUC ≈ -0.03 to -0.04). The larger drops in κ (-0.06 to -0.15, largest for Type A) therefore appear to reflect miscalibration of the default 0.5 threshold under prompt shift rather than a comparable loss of ranking ability. This argues for item-type-specific decision thresholds at deployment.

The third finding is that a single pooled detector is plausible in the familiar-prompt regime. Specialist advantages are small and sit within the pooled model's uncertainty intervals. Under cross-prompt evaluation, the pooled model also remains close to the specialists, but all item-type slices show some loss relative to specialist models, leaving a stronger case for specialization in that regime (Table 7, Table A1).

The main interpretive constraint is still the label-generation process. As established in §1.2, the reported metrics estimate agreement with the existing review process rather than construct-level accuracy. Reference-conditioned recall and unconfirmed flags make this blind spot visible, but they cannot determine how many model-positive, label-negative cases are genuine missed gaming. That requires new blind review data.

A targeted blind re-review is the most direct remedy. Rather than relabeling the corpus, a small stratified sample sourced from the model labels could be reviewed without model scores or prior labels, with two independent reviewers per response.

Oversampling the model-positive/label-negative cell would estimate the rate of genuine gaming among possible misses, while sampling true negatives and true positives would allow reweighting to the population. A sample in the low hundreds would be sufficient to distinguish whether a model's additional flags are mostly noise or include substantial missed gaming, and would also provide

Item type	Regime	Winning configuration	ROC-AUC	κ	Δ
Type A	within-prompt	ensemble: longformer-base · PCA=100 + handcrafted	0.8787 [0.863, 0.893]	0.6130 [0.580, 0.647]	—
Type A	cross-prompt	ensemble: longformer-large · PCA=20 + handcrafted	0.8402 [0.823, 0.856]	0.4650 [0.427, 0.502]	−0.1480
Type B	within-prompt	ensemble: longformer-large · PCA=200 + handcrafted	0.9481 [0.925, 0.968]	0.7487 [0.683, 0.810]	—
Type B	cross-prompt	ensemble: longformer-large · PCA=10 + handcrafted	0.9105 [0.881, 0.938]	0.6548 [0.578, 0.727]	−0.0939
Type C	within-prompt	ensemble: longformer-large · PCA=200 + handcrafted	0.9278 [0.917, 0.938]	0.6950 [0.663, 0.725]	—
Type C	cross-prompt	ensemble: longformer-large · PCA=200 + handcrafted	0.8964 [0.883, 0.910]	0.6310 [0.597, 0.665]	−0.0640
Pooled	within-prompt	ensemble: longformer-large · PCA=200 + handcrafted	0.8986 [0.883, 0.913]	0.6330 [0.595, 0.668]	—
Pooled	cross-prompt	ensemble: longformer-large PCA=50 + handcrafted	0.8597 [0.843, 0.876]	0.5590 [0.521, 0.596]	−0.0740

Table 4: Best-performing configuration per item type and prompt regime, selected by ROC-AUC, with κ and delta κ also reported. Bracketed values indicate 95% percentile bootstrap confidence intervals.

the reliability estimate the present corpus lacks.

Limitations

The study has three main limitations. First, no blind double-read is available, so we have no inter-rater reliability information, and the reference standard is not independent of the detector that helps produce it (§1.2, §1.3).

Second, recorded negatives may contain undetected positives, so precision and recall are relative to observed labels and should not be read as construct-level accuracy.

Third, cross-prompt uncertainty is understated by response-level bootstrap intervals because prompt-level variance is not fully captured, especially for item types with few held-out prompts.

Conclusion

Transformer representations improve gaming detection over handcrafted features across most item types and both generalization regimes. A single pooled model may be defensible in the familiar-prompt regime, while frozen embeddings provide no significant advantage in the most data-constrained unseen-prompt condition. These findings support transformer-aided gaming detection, but only as estimates of agreement with review process labels.

The main constraint is the label-generation process. Review process labels identify gaming that was detected; they do not reveal how much gaming was missed. Positive–unlabeled class-prior estimation can sometimes estimate the share of hidden positives among unlabeled cases (Elkan and Noto, 2008; Bekker and Davis, 2020), but this requires assumptions that are unlikely to hold when the reference detector catches the most conspicuous cases. The missed-positive rate therefore cannot be recovered from observed labels alone.

Future studies should report how labels were produced, whether reviewers saw model output, whether review was routed by a reference detector, and whether any blind double-scoring was available. When such evidence is absent, reliability should be described as un-estimated rather than low. When feasible, targeted blind re-review of model-positive/label-negative cases should be added.

References

Isaac I. Bejar, Michael Flor, Yoko Futagi, and Chaitanya Ramineni. 2014. On the vulnerability of automated scoring to construct-irrelevant response strategies (cirs): An illustration. *Assessing Writing*, 22:48–59.

Isaac I. Bejar, Waverely Hester VanWinkle, Nitin Madnani, William Lewis, and Michael Steier. 2013. Length of textual response as a construct-irrelevant

Regime	Type	Baseline	Extended	Best frozen embedding	Ensemble
within_prompt	Type A	0.7421 [0.722, 0.764]	0.8052 [0.787, 0.825]	0.8721 [0.856, 0.887]	**0.8787 [0.863, 0.893]**
within_prompt	Type B	0.8468 [0.806, 0.886]	0.8962 [0.864, 0.926]	**0.9491 [0.929, 0.968]**	0.9481 [0.925, 0.968]
within_prompt	Type C	0.7057 [0.683, 0.728]	0.7757 [0.756, 0.795]	0.9266 [0.916, 0.937]	**0.9278 [0.917, 0.938]**
within_prompt	all	0.7431 [0.720, 0.765]	0.7781 [0.756, 0.799]	0.8894 [0.873, 0.905]	**0.8986 [0.883, 0.913]**
cross_prompt	Type A	0.7151 [0.692, 0.736]	0.7960 [0.775, 0.816]	0.8125 [0.794, 0.830]	**0.8402 [0.823, 0.856]**
cross_prompt	Type B	0.8589 [0.821, 0.894]	0.9100 [0.880, 0.937]	0.7914 [0.748, 0.832]	**0.9105 [0.881, 0.938]**
cross_prompt	Type C	0.7230 [0.701, 0.746]	0.7856 [0.766, 0.805]	0.8910 [0.877, 0.905]	**0.8964 [0.883, 0.910]**
cross_prompt	all	0.7467 [0.724, 0.769]	0.7786 [0.757, 0.800]	0.8463 [0.828, 0.863]	**0.8597 [0.843, 0.876]**

Table 5: Test-set ROC-AUC for handcrafted, frozen-embedding, and ensemble configurations, per item type and regime. Highest ROC-AUC values are bolded per row. Within each family, the reported configuration is the one selected by highest nested cross-validation ROC-AUC on training data.

Regime	Type	Recall (reference detector caught)	Recall (reference detector missed)	Unconfirmed flags (traditional false positives)
within_prompt	Type A	0.9487 (n=117)	0.7754 (n=846)	184
within_prompt	Type B	0.9355 (n=62)	0.7752 (n=129)	15
within_prompt	Type C	0.9684 (n=95)	0.8420 (n=861)	156
cross_prompt	Type A	0.9505 (n=101)	0.5672 (n=841)	141
cross_prompt	Type B	0.9881 (n=84)	0.6514 (n=109)	28
cross_prompt	Type C	0.9579 (n=95)	0.8237 (n=868)	205

Table 6: Recall on positives the reference detector flagged versus positives it missed, with unconfirmed flag counts. Higher recall on reference-caught than reference-missed positives indicates that the model partly reproduces reference-detector signal, while nontrivial missed-cell recall captures the potential added coverage.

response strategy: The case of shell language. *ETS Research Report Series*, 2013(1):i–39.

Jessa Bekker and Jesse Davis. 2020. Learning from positive and unlabeled data: A survey. *Machine Learning*, 109(4):719–760.

Aoife Cahill, Martin Chodorow, and Michael Flor. 2018. Developing an e-rater advisory to detect babel-generated essays. *Journal of Writing Analytics*, 2(1):203–224.

Gavin C. Cawley and Nicola L. C. Talbot. 2010. On over-fitting in model selection and subsequent selection bias in performance evaluation. *Journal of Machine Learning Research*, 11:2079–2107.

Yuning Ding, Brian Riordan, Andrea Horbach, Aoife Cahill, and Torsten Zesch. 2020. Don't take "nswvtnvakgxpm" for an answer: The surprising vulnerability of automatic content scoring systems to adversarial input. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 882–892.

Charles Elkan and Keith Noto. 2008. Learning classifiers from only positive and unlabeled data. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 213–220.

Haibo He and Edwardo A. Garcia. 2009. Learning from imbalanced data. *IEEE Transactions on Knowledge and Data Engineering*, 21(9):1263–1284.

Derrick Higgins and Michael Heilman. 2014. Managing what we can measure: Quantifying the susceptibility of automated scoring systems to gaming behavior. *Educational Measurement: Issues and Practice*, 33(3):36–46.

Sarah R. Hughes and Rose Clesham. 2024. Maintaining test integrity through integration of human judgement and ai systems. Technical report, Pearson.

Anubha Kabra, Mehar Bhatia, Yaman Kumar Singla, Junyi Jessy Li, and Rajiv Ratn Shah. 2020. Calling out bluff: Attacking the robustness of automatic scor-

Regime	Type	n	Pooled model (slice)	Specialist (own test)	Δ
within_prompt	Type A	666	0.5375 [0.474, 0.604]	0.6130	−0.0755
within_prompt	Type B	390	0.7487 [0.677, 0.806]	0.7487	—
within_prompt	Type C	666	0.6607 [0.604, 0.718]	0.6950	−0.0343
cross_prompt	Type A	666	0.4565 [0.388, 0.523]	0.4650	−0.0085
cross_prompt	Type B	414	0.6280 [0.555, 0.700]	0.6548	−0.0268
cross_prompt	Type C	666	0.6186 [0.556, 0.674]	0.6310	−0.0124

Table 7: Pooled model vs. specialist model Cohen's kappa scores by prompt type and regime. The pooled model is scored within each item-type slice of its own test set; each specialist is scored on its own test set. Bracketed figures show the 95% percentile bootstrap intervals (2,000 resamples) on the pooled slice; specialist figures are point estimates. Here, pooled slices do not easily allow for ROC-AUC to be computed, so κ is reported.

ing systems with simple adversarial testing. *arXiv preprint arXiv:2007.06796*.

Susan Lottridge, B. Godek, A. Jafari, and M. Patel. 2020. Comparing the robustness of deep learning and classical automated scoring approaches to gaming strategies. White paper. https://files.portal.cambiumast.com/corporate-site/documents/CAI-Cambium-Comparing-Robustness-Automated-Scoring-Approaches.pdf.

Robert J. Mislevy, Linda S. Steinberg, and Russell G. Almond. 2003. On the structure of educational assessments. *Measurement: Interdisciplinary Research and Perspectives*, 1(1):3–62.

Les Perelman. 2020. The babel generator and e-rater: 21st-century writing constructs and automated essay scoring (aes). *Journal of Writing Assessment*, 13(1).

Les Perelman, Louis Sobel, Milo Beckman, and Damien Jiang. 2014. Basic automatic b.s. essay language generator (babel). Software.

Yashad Samant, Lee Becker, Scott Hellman, Bradley Behan, Sarah Hughes, and Joshua Southerland. 2025. Automatic detection of inauthentic templated responses in english language assessments. *arXiv preprint arXiv:2509.08355*.

Yaman Kumar Singla, Swapnil Parekh, Somesh Singh, Junyi Jessy Li, Rajiv Ratn Shah, and Changyou Chen. 2023. Aes systems are both overstable and oversensitive: Explaining why and proposing defenses. *Dialogue & Discourse*, 14(1):1–33.

David M. Williamson, Xiaoming Xi, and F. Jay Breyer. 2012. A framework for evaluation and use of automated scoring. *Educational Measurement: Issues and Practice*, 31(1):2–13.

A Supplementary Results

Regime	Type	Metric	Pooled model (slice) [95% CI]	Specialist (own test)
within_prompt	Type A	kappa	0.538 [0.474, 0.604]	0.613
		F1	0.766 [0.730, 0.802]	0.805
		precision	0.775 [0.730, 0.820]	0.812
		recall	0.757 [0.711, 0.804]	0.797
		specificity	0.781 [0.737, 0.825]	0.816
within_prompt	Type B	kappa	0.749 [0.677, 0.806]	0.749
		F1	0.871 [0.831, 0.903]	0.868
		precision	0.897 [0.850, 0.939]	0.915
		recall	0.846 [0.795, 0.893]	0.826
		specificity	0.903 [0.858, 0.942]	0.923
within_prompt	Type C	kappa	0.661 [0.604, 0.718]	0.695
		F1	0.832 [0.800, 0.861]	0.848
		precision	0.824 [0.783, 0.863]	0.845
		recall	0.841 [0.801, 0.879]	0.851
		specificity	0.820 [0.778, 0.859]	0.844
cross_prompt	Type A	kappa	0.456 [0.388, 0.523]	0.465
		F1	0.746 [0.710, 0.780]	0.694
		precision	0.700 [0.655, 0.746]	0.811
		recall	0.799 [0.757, 0.843]	0.606
		specificity	0.658 [0.603, 0.707]	0.859
cross_prompt	Type B	kappa	0.628 [0.555, 0.700]	0.655
		F1	0.812 [0.770, 0.852]	0.822
		precision	0.822 [0.769, 0.874]	0.849
		recall	0.802 [0.746, 0.854]	0.797
		specificity	0.826 [0.775, 0.876]	0.858
cross_prompt	Type C	kappa	0.619 [0.556, 0.674]	0.631
		F1	0.816 [0.782, 0.846]	0.819
		precision	0.789 [0.745, 0.830]	0.803
		recall	0.844 [0.803, 0.882]	0.836
		specificity	0.775 [0.729, 0.818]	0.795

Table A1: Pooled model versus specialist (own-test) performance for all agreement metrics, by item type and prompt regime (both within- and cross-prompt), with 95% percentile bootstrap confidence intervals (2,000 resamples) on the pooled slice; specialist figures are point estimates. ROC-AUC is omitted because it is not computed on the pooled item-type slices; threshold-free ranking quality is reported per split in Table 4.

Split	Winner	vs. handcrafted	vs. frozen embedding
within_prompt/Type A	ensemble_longformer-base_PCA=100	p=1.6e-12 (270/128) **	p=0.187 (82/65)
within_prompt/Type B	ensemble_longformer-large_PCA=200	p=0.016 (34/16) *	p=0.472 (13/18)
within_prompt/Type C	ensemble_longformer-large_PCA=200	p <0.001 (411/112) **	p=0.775 (26/23)
within_prompt/all	ensemble_longformer-large_PCA=200	p <0.001 (295/92) **	p=0.948 (120/118)
cross_prompt/Type A	ensemble_longformer-large_PCA=20	p=0.293 (296/270)	p=5.2e-13 (121/31) **
cross_prompt/Type B	ensemble_longformer-large_PCA=10	p=0.648 (11/8)	p=1.5e-06 (85/32) **
cross_prompt/Type C	ensemble_longformer-large_PCA=200	p <0.001 (356/152) **	p=0.178 (33/22)
cross_prompt/all	ensemble_longformer-large_PCA=50	p=1.7e-11 (269/133) **	p=0.002 (183/128) **

Table A2: McNemar paired comparisons between the winning configuration and each rival model family, by split. Parenthetical values give discordant-pair counts in the direction winner-only/rival-only. * p <0.05; ** p <0.01.

Calibrated, Interpretable Automated Scoring with LLM Likelihoods

S. Thomas Christie
Renaissance Philanthropy
thomas.christie@renphil.org

Markus Hauru
Renaissance Philanthropy
markus.hauru@renphil.org

Anna N. Rafferty
Carleton College
arafferty@carleton.edu

Abstract

Some learning is difficult to capture using selected-response items and is better revealed through tasks involving writing and other rich behavior. Automated scoring of such work products has proven difficult, and typical scoring methods do not support principled degree-of-belief quantification for individual scores or evidence accumulation across tasks. We investigate whether generative models can be used to quantify the evidence provided by student writing. In particular, we use conditional token likelihoods produced by LLMs to score student answers in the SciEntsBank dataset. We find that, for predicting expert labels, zero-shot likelihood-based methods provide scores with accuracy comparable to, and calibration much better than, LLM 'judges'. Well-calibrated scores support response-level and student-level uncertainty quantification, enabling more trustworthy, useful, and interpretable automated scoring.

1 Introduction

In educational assessment, scoring processes often reduce rich cognitive and behavioral outputs to a single value: a half-page of mathematical notation might be coded as 'incorrect', or an explanation of a scientific phenomenon reduced to a 4/5. Scores summarize the adherence of a final work product to a reference or rubric but rarely record the scorer's uncertainty about whether the assigned score is correct. In effect, scores are treated as outputs of an infallible classifier.

A score should reflect a claim about a student's performance on a particular task (Mislevy et al., 2003; DeCarlo et al., 2011), and should express the *degree of evidence* that the observed behavior supplies for the claim. To achieve this, we need a scoring mechanism that produces a probability distribution over claims, rather than categorical labels. Well-calibrated probabilities support evidence accumulation, determinations of required assessment lengths, and decisions about which responses warrant human review. Because probabilities reflect small accumulations of evidence that categorical labels mask, computing them incrementally can even indicate which components of a student's response provide evidence for a claim, a critical signal for interpretability and bias analysis.

In this paper, we evaluate and compare the performance of three scoring methods powered by large language models (LLMs) on responses in the SciEntsBank dataset. We evaluate scoring performance along three axes: accuracy at reproducing human labels, prediction calibration, and prediction interpretability. While classification accuracy is important, the main focus of our analysis is confidence calibration. A well-calibrated scorer is one whose confidence in its score is an accurate measure of the probability that the score is correct. While typical LLM-powered scoring methods use LLMs to generate text labels for scores, we show that two methods that leverage LLM token likelihoods achieve much better calibration than a judge approach while maintaining similar accuracy.

2 Background

In recent work, we argued that because LLMs and other generative methods compute conditional likelihoods natively, they may be able to operate as measurement equations for student-produced text (Christie et al., 2026). If so, student skills and abilities may be directly inferable from rich, naturalistic behavior using probabilistic machinery. In the proposed approach, a student's attribute is represented in natural language as a prompt component Y, with various attribute values (say, 'proficient' versus 'not proficient') encoded as $Y_1, \ldots, Y_n$. An LLM M is used to compute, for each attribute value, the likelihood of the student output X given the item description D:

$$p_M(X \mid Y_i, D) \text{ for all } i = 1, \ldots, n$$

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 168–176
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

From these values, a posterior distribution over the mutually exclusive attribute values $Y_1, \ldots, Y_n$ can be computed. This method draws on an approach called the *noisy channel model* (henceforth *channel*; Min et al., 2022), and has been explored for text classification but had not previously been applied to educational measurement. The method is named after a formalism in information theory in which a message is transmitted over a noisy channel and observed by the receiver, who is tasked with inferring the most likely source message.

In a similar approach, the LLM likelihood model can be used to compute the conditional probabilities of a label directly, namely,

$$p_M(Y_i \mid X, D) \text{ for all } i = 1, \ldots, n$$

which are then normalized. Following Min et al. (2022), we refer to this as the 'direct' method. Like the channel method, the direct method derives a posterior over labels from LLM token probabilities.

In the common 'LLM-as-a-judge' paradigm (henceforth 'judge'; Gu et al., 2026), a model is prompted with a rubric and a student-written text and asked to generate a score or label. While this approach can produce strong classification results (e.g., Naismith et al., 2023), scores are not accompanied by principled estimates of score uncertainty. Several methods for deriving confidence estimates from LLM judgments have been proposed, including simply asking the LLM to state its confidence alongside its score (Tian et al., 2023). These uncertainty estimates are themselves generated text rather than readouts of the model's conditional likelihoods; they reflect a model's self-reported belief, which may not be consistent with its underlying probabilities.

Automated scoring methods can be broadly categorized as supervised (informed by many examples) or zero-shot. Supervised scoring methods can achieve high accuracy by training specialized models (e.g., Kaya and Cicekli, 2024) or fine-tuning LLMs (Karim et al., 2025). While the literature reporting calibration is limited, some supervised approaches also achieve strong calibration (Kapoor et al., 2024; Lade et al., 2026). Despite this, we focus on zero-shot scoring methods for two reasons. First, supervised methods require substantial expert-labeled training data, which is often costly or unavailable. Second, we envision a future in which students can flexibly choose how to demonstrate their learning, and downstream scoring and measurement are adaptable enough to produce valid and reliable inferences regardless. We view the current work as a step toward that future.

3 Data

We evaluate channel, direct, and judge methods on the SciEntsBank dataset (Dzikovska et al., 2013), which consists of 3rd–6th grade student responses to short-answer science questions. We limit analysis to the "unseen answers" (`test_ua`) portion of the dataset, which has 540 responses. The human-annotated ground truth differentiates among types of incorrect answers. Following the 2-way classification task (Dzikovska et al., 2013), we map the five labels to correct-versus-incorrect, turning the problem into a binary classification task.

4 Methods

Consider an item description D, a student's answer to it X, and labels Y_C and Y_I indicating whether the student's answer is judged correct or incorrect, respectively. D may include other context, such as a reference answer or a scoring rubric, depending on the data. In our experiments, D is composed of the item question, the reference answer, and brief descriptions of the two labels. We wish to compute the posterior probabilities of the labels given the student's answer: $p(Y_i \mid X, D)$. We use three different methods to elicit these score probabilities from an LLM and compare the resulting predictions to human labels. Our experiments focus on small and mid-sized open-weight LLMs as base models; the judge can alternatively use a closed-weight model, so we include a comparison to OpenAI's GPT models for reference.

4.1 Score elicitation approaches

Of our three score elicitation approaches, the two likelihood-based methods require an LLM to produce the probability of each possible next token (e.g., a word or word piece), given the preceding text, rather than generate new text. We denote these next-token probabilities by p_M for a model M.

The probability assigned by M to a longer text made of tokens t_0 to t_N, denoted $p_M(t_{\leq N})$, is the product of the probabilities assigned to each consecutive token, conditioned on all preceding tokens:

$$p_M(t_{\leq N}) = p_M(t_0) \prod_{j=1}^{N} p_M(t_j \mid t_{\leq j-1})$$

Property	Judge	Direct	Channel
Native per-token evidence	×	×	✓
Immune to surface-form competition	×	×	✓
Runs over closed (API-only) models	✓	×	×
Supports "thinking" models	✓	×	×

Table 1: Qualitative comparison of the three uncertainty-aware LLM-based scoring methods we study. ✓/× indicate whether a method provides the property.

Access to next-token probabilities typically requires an open-weight LLM, since most closed-model APIs do not expose next-token probabilities for all tokens.

Direct method: In the direct method, we prepare a text that includes the item description D and student answer X. We use the LLM to compute $p_M(Y_i \mid X, D)$ for $i \in \{C, I\}$ and normalize the values, treating the result as the posterior distribution over labels. Some scaffolding is required to make the text read like natural language for the LLM; an abbreviated example prompt is:

> A student was asked "How can one reduce the current in a circuit?" The student's answer was "by increasing the resistance". The student's answer is

The LLM M is queried for the probabilities that the next token in this text is 'correct' or 'incorrect'.

To mitigate the bias whereby an LLM may favor one label or the other regardless of the student's answer, we apply the content-free calibration from Zhao et al. (2021). In particular, we apply this calibration question-by-question, by using a prompt that includes the question and the reference answer, and that replaces the student answer with the string `[REDACTED]`; see Zhao et al. (2021) for details.

Noisy channel model: In the noisy channel model, the order of the student answer and label is inverted. The prepared prompt includes D and Y_i, and we compute the probability $p_M(X \mid Y_i, D)$. As noted above, this is structurally similar to a psychometric measurement equation. We compute the posterior probability using Bayes' theorem:

$$p(Y_i \mid X, D) = \frac{p_M(X \mid Y_i, D)\, p(Y_i \mid D)}{\sum_{j \in \{I,C\}} p_M(X \mid Y_j, D) p(Y_j \mid D)}.$$

Here $p(Y_i \mid D)$ is the prior over answer correctness, which we set to be uniform. The abbreviated example prompt for the channel method is:

> A student was asked "How can one reduce the current in a circuit?" The student's answer was correct. The student's answer was

and we use M to compute the conditional likelihood of the continuation 'by increasing the resistance'. The procedure is repeated by replacing 'correct' with 'incorrect' in the prompt, and the two likelihoods are normalized.

LLM-as-a-judge: The LLM-as-a-judge (judge) method prompts an LLM for both a label and a verbalized confidence rather than using likelihood readouts directly. With the same abbreviated example as above, the prompt is:

> A student was asked "How can one reduce the current in a circuit?" The student's answer was "by increasing the resistance". Is the student's answer correct or incorrect? Answer with JSON: "label": "<correct or incorrect>", "confidence": <probability between 0.0 and 1.0 that the label is correct>.

Note that the judge method should be used with an instruction-tuned LLM, whereas the direct and channel methods work best without instruction fine-tuning; in our experiments, the judge uses each model's instruction-tuned variant while the channel and direct methods use the corresponding base model.

4.2 Method affordances

Method affordances are summarized in Table 1. The channel method natively provides per-token evidence, since it accumulates its score probabilities by computing and comparing the conditional likelihoods of each subsequent student-written token. Similar per-token results can be obtained for the direct and judge methods using prompts where only a partial student answer is shown (see Figure 4). However, this requires rerunning the scoring

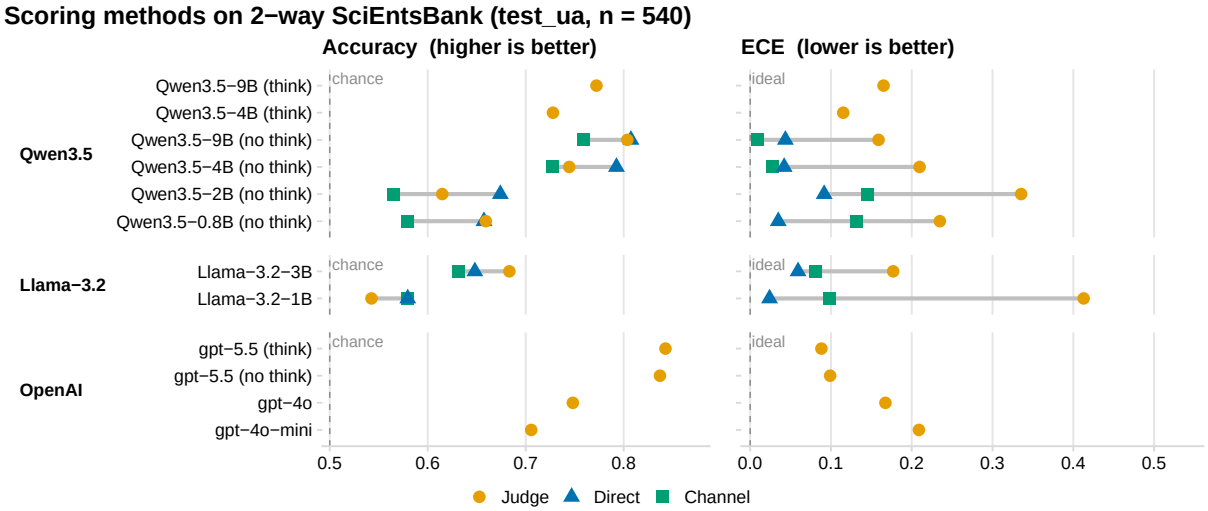


Figure 1: Scoring methods on 2-way SciEntsBank (`test_ua` split, $n = 540$), evaluated with various models. On the left accuracy, on the right 10-bin Expected Calibration Error (ECE). The direct method uses the content-free calibration correction (Zhao et al., 2021). Results typically improve with model size for all methods, and accuracies are broadly similar across them when using the same base LLM. Calibration error, however, is far larger for the LLM-as-judge's self-reported confidence than with the LLM-likelihood-based approaches. The frontier-model judge with GPT-5.5 is superior in accuracy, though still behind the direct method and the channel method with Qwen3.5-9B in calibration error. Our closed-model judge results are broadly in line with those reported by others (Kortemeyer, 2024).

for every token in the answer, whereas the channel method captures this information in one pass.

The channel method is also the only method immune to surface-form competition, a known issue in which probability mass is split across multiple surface forms of the same label, diluting the signal. Methods for mitigating it exist (Holtzman et al., 2021), but are imperfect and complicate implementation.

Both channel and direct methods require open-weight models, so the judge is the only option when only closed-weight models are available.

Finally, LLMs that use thinking traces can, in principle, be used with the direct and channel methods, but are a more natural fit for LLM-as-a-judge. We have not found a way to insert thinking traces into the channel and direct methods without biasing their confidence estimates.

5 Results

Score-prediction results for SciEntsBank are shown in Figure 1.[1] All three scoring methods exhibit comparable accuracy for size-matched models, with the channel method performing slightly worse than the other two. The base LLM affects accuracy more than the scoring method. The small, open models we evaluated are competitive with GPT-4o in accuracy but are outperformed by GPT-5.5.

To measure calibration of confidence estimates, we use Expected Calibration Error (ECE; Naeini et al., 2015; Guo et al., 2017), defined as $\sum_{k=0}^{K-1} \frac{|B_k|}{N} |\mathrm{acc}(B_k) - \mathrm{conf}(B_k)|$, where the predictions from the method have been split into K equal-width confidence bins B_k over $[0, 1]$ based on the probability assigned to the top-scoring label, and $\mathrm{acc}(B_k)$ and $\mathrm{conf}(B_k)$ are the average accuracy and confidence within that bin. We use $K = 10$ bins. As shown on the right in Figure 1, in size-matched comparisons, the channel and direct methods exhibit better-calibrated probability estimates than LLM-as-a-judge.

The reliability plots in Figure 2 provide more detail about prediction calibration for each method, using the Qwen3.5-9B model. The channel method uses the full range of possible confidences and demonstrates good calibration: when the model says it is 60% sure of its answer, it is correct approximately 60% of the time. The direct method is similar, if slightly worse. In contrast, the judge uses only a few stereotypical confidence ratings (e.g., 95% and 100%) and is systematically overconfident.

Figure 3 shows coverage and accuracy when

[1] Source code for results is at github.com/RenaissancePhilanthropy/Calibrated-LLM-scoring-AIME-2026.

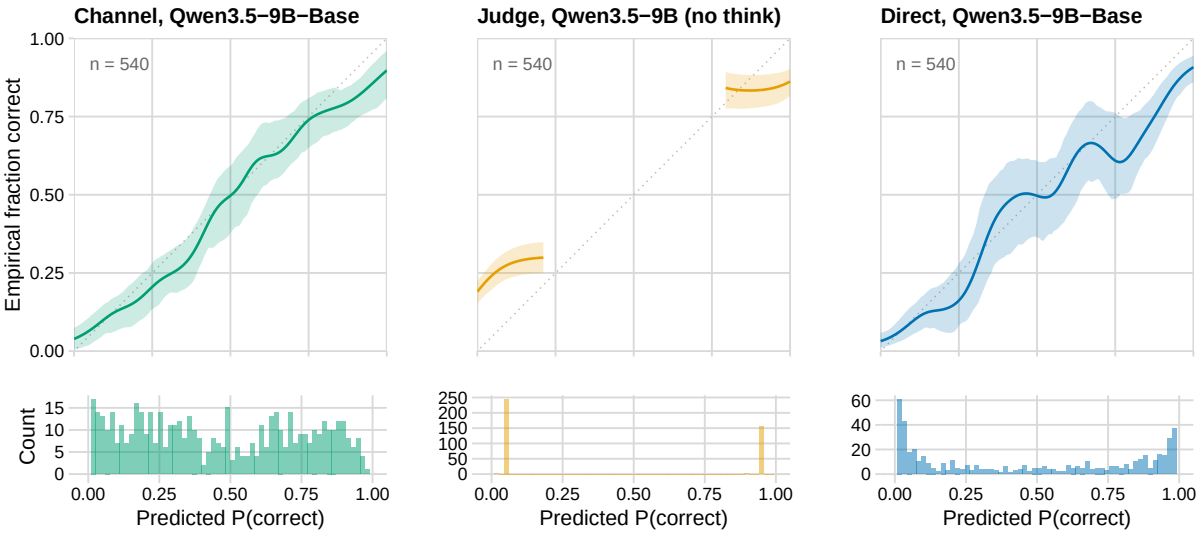


Figure 2: Top: Predicted probability that a student answer is correct versus how often the answer actually was correct. Each sample contributes one point at its predicted $\mathrm{P}(\text{correct})$; the curve is a Gaussian-smoothed ($\sigma = 0.05$) empirical frequency of the gold label being *correct*, with a bootstrapped 95% confidence band, and the dotted diagonal marks perfect calibration. Bottom: The same data, showing its distribution along the predicted $\mathrm{P}(\text{correct})$.

varying a confidence threshold for each of the scoring methods on SciEntsBank, again using Qwen3.5-9B. To produce this plot, for each method we filter to keep only predictions where the model's confidence is above a chosen threshold. We compute the accuracy within this high-confidence subset and the fraction of all predictions it contains (coverage). We then vary the confidence threshold from 0 to 1, and Figure 3 reports, for each required accuracy level, the largest coverage achievable at that accuracy. In practice, a user may choose to trust high-confidence model predictions and route low-confidence scores to a human scorer. While direct and judge methods have roughly equivalent overall accuracy, the better-calibrated direct method allows responses to be ranked by trustworthiness, supporting downstream decisions about when and whether to rely on model outputs.

We explore score interpretability by investigating whether the resulting score can be attributed to specific tokens in the student answers. While a comprehensive analysis is beyond the scope of this work (and an appropriate human-labeled dataset is unavailable for SciEntsBank), we use example item responses as case studies. In Figure 4, we compute the per-token evidence and resulting posterior probability distribution for two question/response pairs. In the channel method, the log-likelihood ratio after token j (equal to the posterior log-odds under a uniform prior) is computed directly over the prefix

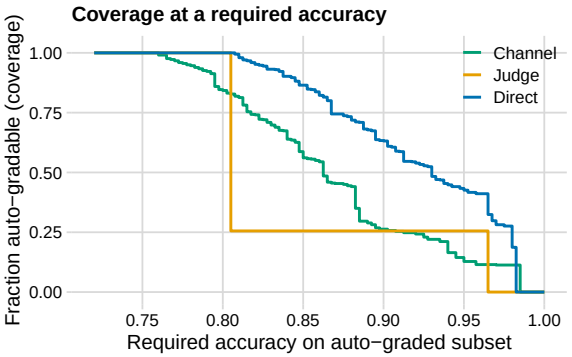


Figure 3: Selective prediction on 2-way SciEntsBank; predictions are from Qwen3.5-9B. For each required accuracy level (horizontal axis), the curve shows the largest fraction of responses that can be auto-scored while meeting that accuracy on the auto-scored subset (vertical axis).

$X_{1\ldots j}$:

$$\mathrm{LLR} = \log\left(\frac{p_M(X_{1\ldots j} \mid Y_C, D)}{p_M(X_{1\ldots j} \mid Y_I, D)}\right)$$

For the direct and judge methods, classifications are made using repeated evaluations of partial (prefix) student responses with the remainder redacted, following the content-free prompt of Zhao et al. (2021). To illustrate, in the first direct-method example, the following prompt would be used to score a partial response:

> A student was asked "How can one reduce the current in a circuit?" The

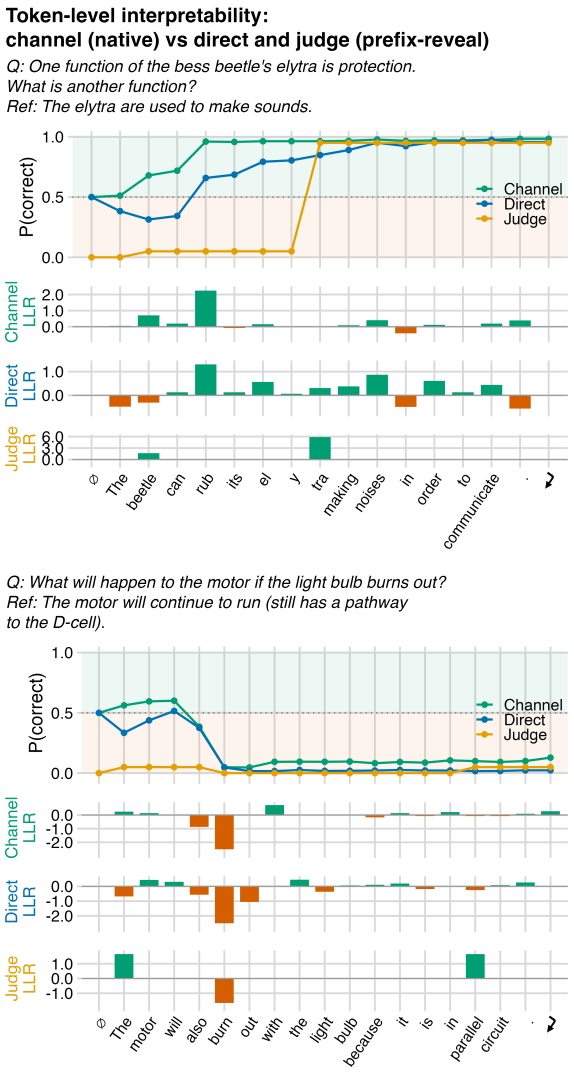


Figure 4: Token-by-token scoring with Qwen3.5-9B on two SciEntsBank question/response pairs. For each example, the top plot overlays each method's cumulative probability that the student answer is correct as the answer is revealed token by token; the dotted line is the 0.5 decision boundary. The three bar plots show each token's log-likelihood-ratio contribution (the change in cumulative LLR), i.e., whether it pushed that method's verdict toward correct (up) or incorrect (down).

> student's answer was "by increasing [REDACTED]". The student's answer is

Both the channel and direct methods show gradual accumulation of evidence throughout the course of the student text, and words with strong probability changes are intuitively indicative of student answer correctness. In contrast, the judge method collects evidence in bursts, a phenomenon that reflects the coarse, clustered confidence values revealed in Figure 2.

6 Discussion

When used in educational assessment, automated scoring is one step in a process of producing valid and reliable estimates of student learning. While most scoring encodes adherence to a rubric, we argue that scores should encode the *degree of evidence* regarding a claim about a student's performance. Probabilistic scoring provides the mathematical mechanism to represent variation in evidence across student responses and even throughout the course of a single response. Moreover, probabilistic scoring places uncertainty directly in the scoring model of an evidence-centered design framework (Mislevy et al., 2003), allowing that uncertainty to be used directly or propagated to other inferences (e.g., those about the student).

We evaluated three zero-shot methods for scoring student writing using LLMs. While the methods showed roughly comparable accuracy across size-matched models, likelihood-based methods produced predictions with consistently superior calibration compared to judges' self-reported confidence. Accuracy and calibration typically improve with model size. While the current study is limited to 2-way classification, other work (Christie et al., 2026, and our preliminary experiments) suggests that these findings likely extend to multi-class settings.

The methods considered here are zero-shot: they require no training data, making them flexible methods for use in assessing performance in novel tasks. However, where human-labeled examples are available, we recommend training a BERT-based classifier or fine-tuning an LLM and using either the channel or direct methods. These methods often produce substantially better accuracy than zero-shot approaches (Ferreira Mello et al., 2025), and in our preliminary experiments we also observed better ECE.

While the LLM judges we evaluated did not provide good uncertainty estimates, this is likely a property of the current generation of LLMs rather than an unavoidable limitation. Recent work demonstrates that penalizing overconfidence during post-training leads to better-calibrated LLM-produced judgments (Damani et al., 2025).

6.1 Limitations

All of the methods we compared are sensitive to prompt wording. Prompt sensitivity suggests a difficult optimization problem, in which tweaking or

rephrasing prompts may improve scoring performance. Automated prompt optimization is possible (Romera-Paredes et al., 2024; Wang et al., 2026), but is beyond the scope of this paper. The choice of bias-correction method also has a dramatic impact on the performance of the direct method. We report the strongest variant, but considered others (Zhao et al., 2021; Fei et al., 2023; Zhou et al., 2024); because this choice was made by comparing variants on the same evaluation set, the direct method's reported calibration may be slightly optimistic.

While we believe that the current work provides a useful case study, we caution readers against overgeneralizing the results. We compared against only a single confidence-elicitation method for LLM judges; others are also available (Farquhar et al., 2024) and may perform differently. We also omitted any comparison to scoring from fine-tuned models or other supervised methods. Early experiments indicated that supervised approaches were substantially better on both accuracy and calibration. The current study evaluates scoring methods on a single dataset, in a single subject area. A much broader analysis is needed to establish the utility of the channel or direct methods for operational settings. Similarly, the number of models compared was small: a more thorough evaluation could include comparisons with larger open-weight models, as well as closed-weight foundation models from different model families.

Finally, while LLM likelihoods provide well-calibrated scores, our investigations also revealed a fundamental limitation: LLM likelihoods score typicality rather than correctness. In other words, a response scores highly because it meets the expectations of the prompted LLM, *not* because it is factually or logically correct. While this is an acceptable or even desirable property for some tasks, it makes likelihood-based scoring an unacceptable method for others. It also indicates a potential source of bias: surface-level dialect or stylistic features underrepresented in a model's training data may unduly affect the scoring process, a phenomenon previously observed when using LLM judges (Ocumpaugh et al., 2025). This challenge is shared by supervised methods, but would nevertheless need to be detected and addressed prior to use with students.

7 Conclusion

Educators and assessors are charged with evaluating student learning against established standards such as the Common Core or state-specific variations (National Governors Association Center for Best Practices and Council of Chief State School Officers, 2010; New York State Education Department, 2017). While selected-response items are a common mechanism for such measurements, students should be able to demonstrate learning in ecologically valid ways, for example through performance tasks or authentic, motivating projects.

Generative models offer the assessment community an opportunity to build scoring and measurement processes that accumulate evidence for proficiency claims from diverse student behavior. A successful method would produce scores flexibly, against multiple claims, from observations of complex student behavior. This would enable dynamic, personalized assessments, while preserving validity and reliability. We suggest that the conditional likelihoods intrinsic to generative modeling provide an underutilized and promising avenue for producing probabilistic judgments about student learning and warrant additional investigation and validation.

Acknowledgments

The authors thank Matthew Zent and Simon Woodhead at Eedi for valuable discussions. The authors also thank the Learning Engineering Virtual Institute, a program of Renaissance Philanthropy, for supporting this work.

References

S. Thomas Christie, Matthew Zent, Markus Hauru, Anna N. Rafferty, and Simon Woodhead. 2026. Likelihood-based diagnosis with generative models: Confidence-aware measurement from student writing. In *Artificial Intelligence in Education: 27th International Conference, AIED 2026, Seoul, South Korea, June 27–July 3, 2026, Proceedings, Part III*, pages 236–244, Berlin, Heidelberg. Springer-Verlag.

Mehul Damani, Isha Puri, Stewart Slocum, Idan Shenfeld, Leshem Choshen, Yoon Kim, and Jacob Andreas. 2025. Beyond binary rewards: Training LMs to reason about their uncertainty. *Preprint*, arXiv:2507.16806.

Lawrence T DeCarlo, YoungKoung Kim, and Matthew S Johnson. 2011. A hierarchical rater model for constructed responses, with a signal detection rater model. *Journal of Educational Measurement*, 48(3):333–356.

Myroslava O Dzikovska, Rodney Nielsen, Chris Brew, Claudia Leacock, Danilo Giampiccolo, Luisa Bentivogli, Peter Clark, Ido Dagan, and Hoa Trang Dang. 2013. SemEval-2013 task 7: The joint student response analysis and 8th recognizing textual entailment challenge. In *Second Joint Conference on Lexical and Computational Semantics (*SEM), Volume 2: Proceedings of the Seventh International Workshop on Semantic Evaluation (SemEval 2013)*, pages 263–274.

Sebastian Farquhar, Jannik Kossen, Lorenz Kuhn, and Yarin Gal. 2024. Detecting hallucinations in large language models using semantic entropy. *Nature*, 630(8017):625–630.

Yu Fei, Yifan Hou, Zeming Chen, and Antoine Bosselut. 2023. Mitigating label biases for in-context learning. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14014–14031, Toronto, Canada. Association for Computational Linguistics.

Rafael Ferreira Mello, Cleon Pereira Junior, Luiz Rodrigues, Filipe Dwan Pereira, Luciano Cabral, Newarney Costa, Geber Ramalho, and Dragan Gasevic. 2025. Automatic short answer grading in the LLM era: Does GPT-4 with prompt engineering beat traditional models? In *Proceedings of the 15th International Learning Analytics and Knowledge Conference*, pages 93–103.

Jiawei Gu, Xuhui Jiang, Zhichao Shi, Hexiang Tan, Xuehao Zhai, Chengjin Xu, Wei Li, Yinghan Shen, Shengjie Ma, Honghao Liu, Saizhuo Wang, Kun Zhang, Zhouchi Lin, Bowen Zhang, Lionel Ni, Wen Gao, Yuanzhuo Wang, and Jian Guo. 2026. A survey on LLM-as-a-judge. *The Innovation*, 7(6):101253.

Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. 2017. On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1321–1330. PMLR.

Ari Holtzman, Peter West, Vered Shwartz, Yejin Choi, and Luke Zettlemoyer. 2021. Surface form competition: Why the highest probability answer isn't always right. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7038–7051.

Sanyam Kapoor, Nate Gruver, Manley Roberts, Arka Pal, Samuel Dooley, Micah Goldblum, and Andrew Wilson. 2024. Calibration-tuning: Teaching large language models to know what they don't know. In *Proceedings of the 1st Workshop on Uncertainty-Aware NLP (UncertaiNLP 2024)*, pages 1–14.

Ahmed Karim, Qiao Wang, and Zheng Yuan. 2025. Beyond the score: Uncertainty-calibrated LLMs for automated essay assessment. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 19642–19647.

Mustafa Kaya and Ilyas Cicekli. 2024. A hybrid approach for automated short answer grading. *IEEE Access*, 12:96332–96341.

Gerd Kortemeyer. 2024. Performance of the pre-trained large language model GPT-4 on automated short answer grading. *Discover Artificial Intelligence*, 4(1):47.

Shirin Lade, Alistair Willis, Jonathan Nylk, and Oli Howson. 2026. Domain-adaptive pre-training for automated short answer grading in conceptual physics: Reliability, question-level analysis, and error reduction. In *Proceedings of the 21st Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2026)*, pages 635–650, San Diego, California, USA. Association for Computational Linguistics.

Sewon Min, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2022. Noisy channel language model prompting for few-shot text classification. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5316–5330.

Robert J Mislevy, Linda S Steinberg, and Russell G Almond. 2003. On the structure of educational assessments. *Measurement (Mahwah NJ)*, 1(1):3–62.

Mahdi Pakdaman Naeini, Gregory F. Cooper, and Milos Hauskrecht. 2015. Obtaining well calibrated probabilities using Bayesian binning. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, AAAI'15, pages 2901–2907. AAAI Press.

Ben Naismith, Phoebe Mulcaire, and Jill Burstein. 2023. Automated evaluation of written discourse coherence using GPT-4. In *Proceedings of the 18th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2023)*, pages 394–403, Toronto, Canada. Association for Computational Linguistics.

National Governors Association Center for Best Practices and Council of Chief State School Officers. 2010. Common core state standards.

New York State Education Department. 2017. New York State P-12 Next Generation Learning Standards for Mathematics and English Language Arts. Technical report, NYSUT Research & Educational Services, Albany, NY.

Jaclyn Ocumpaugh, Xiner Liu, and Andres Felipe Zambrano. 2025. Language models and dialect differences. In *Proceedings of the 15th International Learning Analytics and Knowledge Conference*, LAK '25, pages 204–215, New York, NY, USA. Association for Computing Machinery.

Bernardino Romera-Paredes, Mohammadamin Barekatain, Alexander Novikov, Matej Balog, M. Pawan Kumar, Emilien Dupont, Francisco J. R. Ruiz, Jordan S. Ellenberg, Pengming Wang, Omar Fawzi, Pushmeet Kohli, and Alhussein Fawzi. 2024. Mathematical discoveries from program search with large language models. *Nature*, 625(7995):468–475.

Katherine Tian, Eric Mitchell, Allan Zhou, Archit Sharma, Rafael Rafailov, Huaxiu Yao, Chelsea Finn, and Christopher D Manning. 2023. Just ask for calibration: Strategies for eliciting calibrated confidence scores from language models fine-tuned with human feedback. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5433–5442.

Yun Wang, Xin Xia, Xuansheng Wu, Xiaoming Zhai, and Ninghao Liu. 2026. Learnable assessment skills for LLM-based automated scoring: Rubric construction via iterative optimization. *Preprint*, arXiv:2605.29274.

Zihao Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. 2021. Calibrate before use: Improving few-shot performance of language models. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 12697–12706. PMLR.

Han Zhou, Xingchen Wan, Lev Proleev, Diana Mincu, Jilin Chen, Katherine Heller, and Subhrajit Roy. 2024. Batch calibration: Rethinking calibration for in-context learning and prompt engineering. In *The Twelfth International Conference on Learning Representations (ICLR)*.

Predicting Item-to-Range Performance Level Descriptor Matches with Structured Language Model Reasoning

Kang Xue, M. Christina Schneider, Sangdon Lim, Garron Gianopulos, and Alexandra Perez
Cambium Assessment, Inc.

Abstract

Range performance level descriptors (RPLDs) connect assessment items to claims about what students at different achievement levels know and can do. Retrospectively assigning RPLDs to a large item bank is valuable but labor intensive. This work-in-progress study evaluates whether small and large language models can assist expert item-to-RPLD matching for 524 Grade 3–5 English language arts items. We compared direct classification with structured prompts that guide a model to analyze the knowledge, skills, evidence, and cognitive demand required by an item. We also examined self-consistency voting and a rater-informed prompting. Preliminary results indicate that larger hosted models achieved the closest overall performance to humans, although some locally hosted models produced comparable results. Voting improved prediction reliability but did not consistently increase agreement with human scores, whereas training models with human scores generally improved classification accuracy. Model performance also tended to decline as grade level increased.

1 Background

Principled assessment design (Huff et al., 2025) connects claims about student proficiency to the evidence elicited by assessment tasks. Within this approach, Range performance level descriptors (RPLDs) (Egan et al., 2012) support item development and score interpretations by defining evidence characteristics that distinguish increasingly sophisticated levels of student understanding and cognitive processing. Difficulty drivers and other construct-relevant item features can therefore be used intentionally to engineer items for targeted performance definitions found in the RPLDs.

Content specialists must invest considerable time to document alignments between items and RPLDs, which can create human resource challenges. When a state adopts new standards and rewrites RPLDs, existing items may need to be retroactively realigned to new RPLDs and standards. For computer adaptive testing programs, the resulting workload can involve thousands of items. Content specialists must analyze the item, identify the evidence a student demonstrates when responding correctly, compare that evidence with adjacent descriptors, and resolve ambiguous boundary cases to conduct the alignment process. Language models may reduce this burden, but an apparently plausible classification is insufficient: the classification must be accurate, reliable, and relatively consistent when compared to human judgment.

The present work evaluates language models as decision-support tools for item-to-RPLD matching. It asks four questions:

1. How closely do language-model performance-level classifications agree with resolved scores relative to two human raters in terms of exact and adjacent agreement?

2. Does expert-informed prompting improve agreement with resolved expert performance-level classifications relative to direct classification?

3. Does self-consistency voting improve classification accuracy or agreement with resolved expert performance level classifications?

4. Where do model-expert discrepancies occur, particularly among items with lower human consensus or disputed first-round ratings?

We used an evaluation framework that treats RPLD prediction as an ordinal, expert-referenced task rather than ordinary multiclass classification. In addition to exact agreement, we evaluated adjacent-level agreement, quadratic weighted kappa, directional error, coverage, disagreement by human-consensus status, and computational cost.

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 177–183
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

The current results are preliminary because the experiment matrix remains unbalanced across model families and prompting conditions.

2 Data

The dataset consisted of English Language Arts items from Grades 3–5. Each item record included the stimulus, when applicable, item stem, response options, answer key, relevant standard information, the full set of RPLD level definitions, two initial human RPLD ratings, and a final resolved rating. The final analytic sample comprised 524 items with sufficient item-content text and relevant standard information for LLM-based classification: 195 Grade 3 items, 125 Grade 4 items, and 204 Grade 5 items.

RPLD levels were ordered as Below, Approaching, At, and Above. Each item was initially rated by two content experts. Where agreement was not exact, ratings went through resolution for a final alignment.

3 Method

3.1 Models and Experimental Conditions

We evaluated instruction-tuned models from the Qwen2.5 (Yang et al., 2024), Gemma 2 (Gemma Team et al., 2024), Llama 3.1 and 3.2 (Grattafiori et al., 2024; Meta AI, 2024), and DeepSeek-R1-Distill (DeepSeek-AI et al., 2025) families, with parameter sizes ranging from approximately 1B to 32B. We also evaluated the hosted Claude Sonnet 5 model (Anthropic, 2026), accessed through Amazon Bedrock. All open-weight models were deployed on an NVIDIA L40S 48GB GPU. To accommodate the memory requirements of models of different scales, smaller models were executed using FP16 precision, whereas models with 14B parameters or more were executed using FP8 quantization. For each run, the experimental pipeline recorded the model identifier, prompt condition, sampling configuration, quantization setting, random seed, inference time, and token usage when available.

The voting condition was included to improve prediction consistency and reduce reliance on any single language model response. By aggregating ten independently generated performance level classifications, majority voting can mitigate run-to-run variability and produce a more stable final prediction. This condition was applied to non-DeepSeek model configurations for which repeated inferences were computationally feasible. DeepSeek-R1-Distill models were evaluated only with single-run predictions because their reasoning-oriented generation produced substantially longer inference times than comparably sized instruction-tuned models, making ten-vote self-consistency too computationally time-consuming for the present experiment. The majority voting approach is particularly relevant for models whose temperature (i.e., settings that control the degree of randomness an LLM uses when generating a response) cannot be directly controlled; the repeated inference and aggregation provide an alternative mechanism for assessing and improving output stability (Xue and Appleton, 2026).

We evaluated two prompting conditions. In the direct condition, the model assigned one of four performance levels based on the item and the corresponding standard alignment metadata. In the prior-knowledge condition, the model received content specialist-developed guidance for interpreting item demands and differentiating adjacent RPLD levels. This guidance directed the model to prioritize evidence from the current item and its descriptors, evaluate cognitive demand, information scope, and option plausibility, and examine the boundaries surrounding the best-fitting level. Worked examples were provided solely for calibration rather than as answer keys. In both conditions, the model returned a single normalized RPLD label; missing or unparseable responses were treated as coverage failures.

For local models, deterministic single-run predictions were generated with a temperature of 0. In the voting condition, ten sampled predictions were aggregated by majority vote, with temperature set to 0.3 for models smaller than 8B parameters and 0.6 for models with 8B parameters or more. These settings were selected to introduce sufficient response diversity for self-consistency voting while limiting excessive variation, particularly among smaller models. Because temperature was not configurable for the Amazon Bedrock implementation used in this study, Claude Sonnet 5 was evaluated under the provider-defined low- and high-effort settings. These settings control the amount of computational effort allocated to reasoning: the low-effort setting favors faster, less resource-intensive inferences, whereas the high-effort setting permits more extensive reasoning at greater computational cost.

Model family	Model variants	Parameter size	Deployment	Precision or quantization
Qwen 2.5 Instruct	1.5B, 3B, 7B, 14B, 32B	1.5–32B	L40S 48GB	FP16 or INT8
Gemma 2 Instruct	2B, 9B, 27B	2–27B	L40S 48GB	FP16 or BF16
Llama Instruct	Llama 3.2 1B and 3B; Llama 3.1 8B	1–8B	L40S 48GB	FP16 or BF16
DeepSeek R1 Distill	Qwen 1.5B, 7B, 14B, and 32B; Llama 8B	1.5–32B	L40S 48GB	FP16 or INT8
Claude Sonnet 5	Low- and high-effort settings	Not disclosed	Amazon Bedrock	

Table 1: Models included in the experiment.

3.2 Evaluation

Model performance was evaluated against the Round 2 resolved human classification using four primary metrics: exact agreement, exact plus adjacent agreement, Pearson correlation, and quadratic weighted kappa. Exact agreement represents the proportion of model predictions that matched the resolved performance level, whereas exact plus adjacent agreement represents the proportion that either matched the resolved level or differed by no more than one level. Pearson correlations were used to assess the strength of the linear association between model-assigned and resolved performance levels. The same metrics were calculated for first-round human ratings to provide a descriptive benchmark for contextualizing model performance.

Grade	Group	Configuration	Valid n	Exact	Adjacent	Pearson r	QWK
G3	Human R1	Human rating	145	.876	.979	.914	.910
G3	Human R2	Human rating	100	.850	.960	.877	.869
G3	Human R3	Human rating	145	.917	.972	.900	.896
G3	Claude Sonnet 5	Prior knowledge, voting	192	.708	.917	.662	.661
G3	DeepSeek R1 Distill	Qwen-32B, prior knowledge	193	.637	.870	.547	.528
G3	Qwen 2.5 Instruct	14B, prior knowledge, voting	193	.622	.845	.439	.387
G3	Gemma 2 Instruct	27B, direct	195	.590	.856	.407	.309
G3	Llama Instruct	3.1-8B, voting	190	.568	.805	.341	.214
Grade 4							
G4	Human R1	Human rating	83	.940	.976	.906	.905
G4	Human R2	Human rating	116	.957	.983	.942	.940
G4	Human R3	Human rating	51	.804	.902	.704	.702
G4	Claude Sonnet 5	Prior knowledge, single	125	.736	.880	.557	.556
G4	DeepSeek R1 Distill	Qwen-32B, prior knowledge	125	.504	.856	.411	.406
G4	Qwen 2.5 Instruct	32B, prior knowledge	125	.560	.888	.467	.449
G4	Gemma 2 Instruct	27B, prior knowledge, voting	112	.366	.777	.276	.198
G4	Llama Instruct	3.2-3B, direct, voting	122	.377	.869	-.046	-.037
Grade 5							
G5	Human R1	Human rating	134	.896	.940	.884	.882
G5	Human R2	Human rating	138	.833	.935	.855	.843
G5	Human R3	Human rating	136	.934	.985	.919	.919
G5	Claude Sonnet 5	Prior knowledge, voting	202	.574	.906	.568	.561
G5	DeepSeek R1 Distill	Qwen-32B, prior knowledge	204	.505	.882	.542	.528
G5	Qwen 2.5 Instruct	32B, prior knowledge, voting	203	.448	.911	.464	.424
G5	Gemma 2 Instruct	9B, prior knowledge, voting	203	.399	.833	.113	.098
G5	Llama Instruct	3.2-3B, direct, voting	196	.403	.929	.407	.345

Table 2: Human ratings and best-performing model-family results by grade. QWK denotes quadratic weighted kappa.

4 Results

4.1 Overall Performance Relative to Human Benchmarks

We used the Round 2 resolved ratings as the human benchmark. Across Grades 3–5, individual human raters remained substantially closer to the resolved labels than any model configuration. However, the Round 2 resolved level was established through human review and discussion of the initial Round 1 ratings. Because the resolved level was informed by these initial human judgments, it is not independent of them, which partly explains the high accuracy and correlations observed for the human raters. Table 2 reports all available human-rater comparisons and the best-performing configuration from each model family, organized within each grade. Human exact agreement, adjacent agreement, Pearson correlations, and quadratic weighted kappa values remained higher than all model-family values.

Model performance was meaningfully lower than the human benchmarks, although several configurations produced useful levels of ordinal agreement. The strongest model result in Table 2 was obtained by Claude Sonnet 5. Under the prior-knowledge voting condition, Claude Sonnet 5 reached exact agreement of .708 in Grade 3 and .574 in Grade 5; in Grade 4, its strongest listed result was the prior-knowledge single-run condition, with exact agreement of .736. Adjacent agreement for the Claude configurations was .917, .880, and .906 across Grades 3–5, respectively, indicating that many model errors were within one RPLD level. Among local models evaluated across all three grades, the strongest prior-knowledge results were concentrated among larger models. DeepSeek-R1-Distill-Qwen-32B achieved exact accuracy of .637, .504, and .505 in Grades 3, 4, and 5, respectively, whereas Qwen2.5-32B-Instruct achieved .560 in Grade 4 under prior knowledge and .448 in Grade 5 under prior knowledge with voting. In Grade 3, the best listed Qwen2.5 result was the 14B prior-knowledge voting configuration, with exact accuracy of .622.

Within most model families, the strongest listed configuration tended to involve a larger parameter variant, including DeepSeek-R1-Distill-Qwen-32B and Qwen2.5-32B-Instruct in Grades 4 and 5. The Llama family showed a different pattern in the table: the best listed Llama result was Llama 3.1-8B with voting in Grade 3, whereas Llama 3.2-3B with direct voting (zero-shot voting) was the best listed Llama configuration in Grades 4 and 5. This pattern should be interpreted cautiously because the comparison confounds parameter size with model

generation, prompting condition, and post-training differences; the Grade 4 and Grade 5 advantage of Llama 3.2-3B may therefore reflect differences in calibration to this four-level ordinal classification task rather than a general benefit of smaller model size.

Performance also varied by grade. Grade 3 generally showed stronger model results than Grade 5 for the leading hosted and local configurations: Claude Sonnet 5 declined from .708 exact agreement in Grade 3 to .574 in Grade 5 under prior-knowledge voting, and DeepSeek-R1-Distill-Qwen-32B declined from .637 to .505 under prior knowledge. Qwen2.5 results showed the same general direction when comparing the strongest listed Grade 3 Qwen2.5 configuration (.622 for 14B prior-knowledge voting) with the Grade 5 Qwen2.5-32B prior-knowledge voting result (.448). This pattern suggests that agreement with resolved RPLD performance levels became more difficult in the upper-grade items set represented in the current table.

Adjacent agreement often presented a more favorable picture than exact agreement. Several configurations with modest exact agreement still achieved adjacent agreement at or above .85. However, Pearson correlations and QWK values remained below the human benchmarks and did not meet the criteria typically required for operational deployment in high-stakes automated scoring. For example, the Claude Sonnet 5 configurations produced QWK values of .661, .556, and .561 across Grades 3–5. Nevertheless, these models were developed using a relatively small set of expert-annotated training examples, and the grade-level samples functioned as validation-style evaluation sets. Given the low-stakes context, the observed exact agreement provides preliminary evidence that Claude Sonnet 5 could serve as one initial rater, provided that a human completes the other rating and disagreements are referred for resolution.

4.2 Effect of Self-Consistency Voting

Self-consistency voting did not consistently improve performance across grades, models, or prompting conditions as shown in Figure 1. In most paired comparisons, the differences between single-run and voting predictions were small, and improvements in exact accuracy were often accompanied by minimal gains, no change, or declines in adjacent accuracy, Pearson correlation, or quadratic weighted kappa. For example, under the prior-knowledge condition, Claude Sonnet 5 showed almost no change in Grade 3 exact agreement (.708 to .708) or QWK (.659 to .661), improved exact agreement slightly in Grade 5 (.562 to .574), but decreased in Grade 4 (.734 to .701). Similar mixed patterns appeared for local models: Qwen2.5-32B-Instruct showed small gains in Grade 3 and Grade 5 under prior knowledge, but a slight decline in Grade 4.

These results suggest that voting should not be interpreted as a general accuracy-improving strategy for this task. Instead, its main value is diagnostic: repeated predictions make it possible to evaluate whether a model converges on the same RPLD label across independent generations or whether its classifications are unstable near adjacent descriptor boundaries. When the underlying model already produces stable and relatively unbiased predictions, majority voting may preserve or slightly smooth performance; when the model has a systematic bias, however, voting may simply reinforce the same incorrect label. Therefore, in the present task, self-consistency voting appears more useful for assessing prediction stability and identifying uncertain cases for review than for substantially increasing agreement with resolved expert labels.

4.3 Effect of Prior-Knowledge Prompting

In contrast to self-consistency voting, the prior-knowledge prompt produced more consistent evidence of performance improvement. Across paired direct and prior-knowledge comparisons, adding expert-derived rating guidance generally increased exact agreements, adjacent agreements, Pearson correlations, and quadratic weighted kappas, particularly in Grades 4 and 5 and for larger local models. The largest gains occurred when the direct prompt appeared poorly calibrated to the RPLD scale: for example, Qwen2.5-7B-Instruct in Grade 4 improved from .281 to .469 in exact accuracy under the single-run condition, and Qwen2.5-14B-Instruct in Grade 5 improved from .216 to .343. Similar gains were observed for Claude Sonnet 5, which improved from .664 to .734 in Grade 4 and from .490 to .562 in Grade 5 under single-run comparisons.

Prior-knowledge benefits were not limited to exact agreement. In several cases, the prior-knowledge prompt produced larger gains in Pearson correlation than in exact agreement, suggesting that content specialist-derived guidance helped models better preserve the ordering of the four performance levels described by the RPLDs. For in-

stance, Qwen2.5-32B-Instruct in Grade 5 under voting increased QWK by .231, and Qwen2.5-32B-Instruct in Grade 4 under voting increased QWK by .241. These improvements indicate that the prompt was especially useful for aligning responses with adjacent descriptor boundaries, even when it did not eliminate all exact-agreement errors.

However, the prior-knowledge prompt was not uniformly beneficial. Some Grade 3 configurations declined after adding the expert-derived guidance, most notably Qwen2.5-7B-Instruct, which decreased from .559 to .292 in exact accuracy under the single-run condition, and Llama-3.1-8B-Instruct, which showed a substantial decline in Pearson correlation and QWK. One plausible explanation is that the calibration examples used in the prior-knowledge prompt were drawn from Grade 6 items. Although Grade 3 items were generally simpler, they may occupy a developmental boundary between foundational reading skills and higher-level comprehension or analysis, making Grade 6-based calibration less directly transferable. As a result, the same expert-derived guidance may have aligned better with Grade 4 and Grade 5 items while shifting some Grade 3 models away from the label distribution implied by the item and descriptors. Thus, prior-knowledge prompting appears more effective as a structured task representation than as a guaranteed improvement mechanism.

Overall, these findings support the value of incorporating domain-specific rating knowledge into the prompt, particularly for Grade 4 and Grade 5 items and for models with sufficient capacity to use the additional information. At the same time, the mixed Grade 3 results indicate that prior knowledge should be empirically validated for each model and grade rather than assumed to transfer uniformly across developmental levels. Future analyses should examine whether the negative Grade 3 cases reflect prompt sensitivity, label-distribution shifts, or a mismatch between Grade 6 calibration examples and the foundational comprehension boundary represented in lower-grade RPLD distinctions.

4.4 Consensus Disagreement and Efficiency

Both model and human performance differed substantially between consensus and disputed items. For Claude Sonnet 5 under the prior-knowledge single-run condition, exact agreement declined by 11.3–40.0 percentage points across grades, adjacent agreement declined by 3.6–15.5 percentage points, and Pearson correlations decreased by .386–.715. Human raters also performed less accurately on disputed items, with mean exact agreement ranging from .448 to .463, adjacent agreement from .753 to .869, and Pearson correlations from .348 to .458. Although exact agreement decreased for both humans and the model, relatively high adjacent accuracy indicated that many disagreements remained within one RPLD level of the resolved classification.

These findings suggest that disputed items represent meaningful RPLD boundary cases rather than random disagreement. Future versions of the prior-knowledge prompt should therefore incorporate carefully selected examples of human disagreement and their subsequent resolution to better illustrate ambiguous boundaries and expert decision strategies. To prevent information leakage, these calibration examples should remain separate from the evaluation set.

5 Conclusion and Discussion

The preliminary findings provide qualified evidence that structured, domain-informed prompting can improve item-to-RPLD classification. Improvements were most evident when prompts incorporated expert-developed guidance for interpreting item demands and distinguishing adjacent RPLD levels. This finding suggests that model performance depends not only on model capacity but also on how closely the prompt represents the reasoning required by the assessment framework.

Nevertheless, all model configurations remained below human performance across grades. Claude Sonnet 5, accessed through Amazon Bedrock, achieved the strongest overall model performance, but its agreement with the resolved labels did not reach that of the individual human raters. Although adjacent accuracy was relatively high for several configurations, exact agreement and QWK remained lower than the human benchmarks. These results indicate that identifying a nearby RPLD level is substantially easier for the models than reproducing the specific classification selected through expert review.

From an operational perspective, the observed error patterns support using large language models—particularly Claude Sonnet 5—as initial-rater decision-support tools rather than autonomous RPLD coders. Claude Sonnet 5 achieved approximately 71% exact agreement with the final resolved

ratings, suggesting that models could generate preliminary tags, identify plausible classifications, or flag unstable cases for expert review, thereby concentrating human effort on items that are difficult to classify. In such a workflow, model confidence should not be inferred solely from the fluency of the generated rationale. Agreement across repeated predictions, consistency across prompting conditions, and empirically estimated error rates may provide more defensible indicators for prioritizing human review.

In this study, voting did not consistently improve classification accuracy but provided a useful mechanism for assessing prediction stability and identifying uncertain classifications. For hosted models accessed through online services, where sampling parameters may not be directly configurable, aggregating repeated predictions can reduce reliance on a single response and provide a practical measure of output stability. For locally deployed models, voting involves a trade-off between response diversity and consistency. Selecting an appropriate temperature is therefore important for generating sufficient variation to make voting informative without introducing excessive randomness.

Finally, locally deployed open-weight models may provide a practical compromise between predictive performance and operational security. Although the hosted model achieved the strongest overall results, some larger local models approached its performance, particularly when provided with task-specific prior knowledge. These models captured a meaningful degree of ordinal information while allowing secure assessment content to remain within locally controlled infrastructure. Local deployment may therefore be suitable for pre-tagging, triage, and reviewer prioritization when data security, cost, latency, or infrastructure control are important considerations. However, the observed performance gap indicates that both local and hosted models should remain subject to expert review before their classifications are used operationally.

6 Limitations and Ongoing Work

This is a work-in-progress analysis with several limitations. First, the experiment matrix is unbalanced: some structured-skill runs currently cover only one grade, and not every model appears in every prompt and voting condition. Second, the resolved label is not independent of the original human ratings, so human-to-resolution agreement is a contextual benchmark rather than a direct ceiling. Third, coverage differs across reviewers and some model runs. Fourth, the current results are descriptive and do not yet include paired significance tests or bootstrap confidence intervals. Fifth, findings are limited to Grade 3-5 English language arts items and the RPLDs represented in this item bank. Finally, prompts include operational assessment content and expert-derived guidance; replication will require careful protection of secure material and documentation of prompt versions.

The current benchmark also does not directly quantify the gap between LLM predictions and the two initial human ratings from Round 1, because model performance was evaluated primarily against the Round 2 resolved ratings. Subsequent analyses will therefore compare each model directly with the initial human ratings, using a common item set, to provide a clearer assessment of how model agreement differs from first-round human agreement.

Ongoing work will complete the crossed model-by-prompt-by-voting design on a common item set; estimate paired differences with item-level resampling; analyze confusion matrices and directional bias by grade and RPLD level; characterize disputed items qualitatively; and evaluate selective prediction policies that defer uncertain cases to human reviewers. We will also examine whether lightweight adaptation offers gains beyond prompt engineering, provided that a leakage-resistant train-validation-test split can be constructed.

7 Conclusion

Language models show promise as supplemental tools for retrospective RPLD tagging, particularly when their reasoning is structured around the knowledge, skills, evidence, and boundary distinctions used by expert raters. In the current experiments, a rater-informed hosted model improved over direct prompting but remained well below human agreement. These results favor a human-in-the-loop workflow in which models support prioritization and second review, while experts retain responsibility for final classifications and ambiguous cases.

References

Anthropic. 2026. Introducing claude sonnet 5.

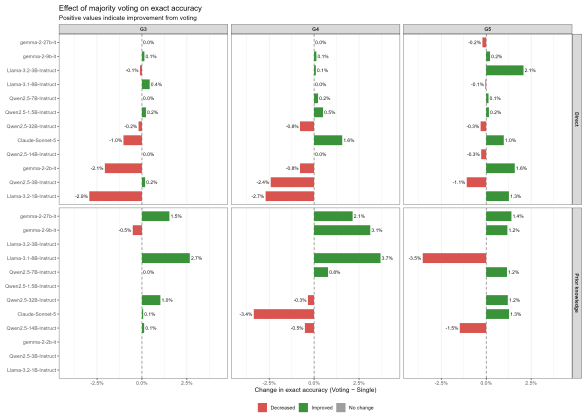


Figure 1: Changes in model performance from single-run prediction to self-consistency voting across grades and evaluation metrics.

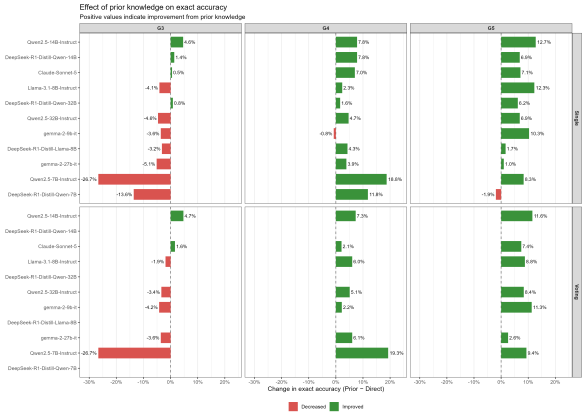


Figure 2: Changes in model performance from direct prompting to prior-knowledge prompting across grades and evaluation metrics.

DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, et al. 2025. DeepSeek-R1: Incentivizing reasoning capability in LLMs via reinforcement learning. *arXiv preprint arXiv:2501.12948*.

Katherine L. Egan, M. Christina Schneider, and Steve Ferrara. 2012. Performance level descriptors. In Gregory J. Cizek, editor, *Setting Performance Standards: Foundations, Methods, and Innovations*, 2 edition, pages 79–106. Routledge.

Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, et al. 2024. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*.

Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, et al. 2024. The Llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.

Kristen Huff, Paul Nichols, and M. Christina Schneider. 2025. Designing and developing educational assessments. In Linda L. Cook and Mary J. Pitoniak, editors, *Educational Measurement*, 5 edition, pages 441–511. Oxford University Press.

Meta AI. 2024. Llama 3.2: Revolutionizing edge AI and vision with open, customizable models.

Kang Xue and James J. Appleton. 2026. Evaluating general-purpose multimodal AI for Q-matrix generation from math items: A cognitive diagnostic modeling exploration. *Journal of Educational Measurement*.

An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, et al. 2024. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*.

Using LLM Judges' Paired Comparisons to Estimate Mathematics Item Difficulty

Lanrong Li, Mohammed A. A. Abulela, Matthew Gushta
MetaMetrics, Inc.
{jli, mabulela, mgushta}@metametrics.com

Abstract

Estimating item difficulty without collecting field testing data has been a long sought-after goal in educational measurement. We prompted large language models (LLMs) to compare mathematics items in pairs to estimate their difficulty. Results showed strong correlations between difficulty based on one LLM judge's paired comparisons and empirical item difficulty.

1 Introduction

Estimating item difficulty without collecting field test data has been a long sought-after goal in educational measurement. This line of research is not only motivated by the potential of substantially reducing the cost and time for estimating the difficulty of newly developed items, but also by the needs of personalized learning and adaptive assessments. A variety of approaches have been proposed to estimate item difficulty without using participants' item response data, spanning using content experts' judgments (e.g., Bejar, 1981), applying traditional statistical models (e.g., Rupp, 2011), machine learning (e.g., Cai et al., 2026), or deep learning models (e.g., Huang et al., 2017) to item features, leveraging large language models (LLMs; e.g., He et al., 2024; Maeda, 2025), or combining multiple methods (Razavi & Powers, 2025). However, each approach has its limitations. For example, using content experts' judgments is costly and time consuming, using feature-based approaches requires pre-extracted item features and known item difficulties, while using LLMs as respondents (e.g., He et al., 2024) or direct difficulty estimators (e.g., Razavi & Powers, 2025) has achieved limited success.

Recently, LLMs have been increasingly used as evaluators to make judgments (e.g., Liu et al., 2024; Liusie et al., 2024; Qin et al., 2023) on natural language processing tasks. Paired comparison (Thurstone, 1927), in which two items are presented and compared at a time, has been shown to be more effective than direct estimation with LLMs (e.g., Liusie et al., 2024). Paired comparison data can be analyzed using the Bradley-Terry model (Bradley & Terry, 1952) to obtain quantitative measures of the attribute of interest for all evaluated candidates. In educational measurement, however, few studies used this approach to estimate item difficulty. One exception is the study by Li and Gushta (2026a), which investigated the performance of LLMs in making paired comparisons to estimate English reading item difficulty. Two other relevant studies are Ballon et al. (2025) and Kolesnikova et al. (2026), but neither controlled for model position bias (Liusie et al., 2024) nor used items containing both text and images. This study aims to use the paired comparisons generated by multimodal LLMs to estimate the difficulty of mathematics items that contain text, mathematical symbols, or images.

2 Background

Earlier methods for estimating item difficulty without using participants' item response data include asking subject matter experts to estimate item difficulty directly (e.g., Lorge & Kruglov, 1952; Bejar, 1983) or make comparative judgments (Attali et al., 2014), and applying regression models to extracted item features (Rupp et al., 2001). With the advancement in data science, the use of machine learning, deep learning, and natural language processing in item difficulty prediction has been increasing (see the review by Alkhuzaey et al., 2024). The emergence of LLMs attracted a

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 184–192
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

surge of interest in their potential to predict item difficulty. Some researchers prompted them to predict item difficulty directly (e.g., Kolesnikova et al., 2026), whereas others used them as respondents to generate item responses for subsequent item calibration (e.g., Maeda, 2025). In these studies, empirical item difficulties based on human participants' item response data were often used as the ground truth. The correlations between predicted and empirical item difficulty observed across approaches ranged from weak to strong. For example, Razavi and Powers (2025) reported correlations ranging from .13 to .55 across individual grade levels and a correlation of .83 when using all 2,564 mathematics items developed for Grades K–5 students.

LLMs, due to their remarkable capability of understanding human languages and performing tasks that require reasoning, a new paradigm of LLM-as-a-judge has emerged (Zheng et al., 2023), in which LLMs are used as judges to assess the quality of the outputs of other AI models or systems. Two methods are often used when employing LLMs as judges: one is to instruct the LLM to provide a rating directly based on predefined criteria, and the other is to instruct it to make pairwise comparisons. Pairwise comparison is a method first proposed by Thurstone (1927) in psychology and was later applied in many other fields. It is simple, intuitive, easy to implement. Paired comparison ratings can be analyzed by the Bradley-Terry model (Bradley & Terry, 1952) to generate estimates of target attributes on an interval scale. Although as the number of candidates to be compared increases, the number of all possible comparisons increases quadratically, not all possible pairs need to be used. Jones and Davies (2023) suggest that 10 comparisons per candidate is sufficient for most cases. Li and Gushta (2026a) has shown that as few as 10 comparisons per item can be used to obtain stable and useful item difficulty estimates. The paired comparison method has also been shown to be more effective than the direct rating method (e.g., Liu et al., 2024; Liusie et al., 2024; Qin et al., 2023).

While using LLMs as judges, their various biases have also been revealed (e.g., Wang, 2023). For example, it has been found that LLMs show position bias during pairwise comparisons (e.g., Liusie et al., 2024), in which they tend to select the first or the second item in a pair more frequently regardless of the amount of the attributes of the two items. The swapping method (Zheng et al., 2023), in which the presentation order of the two items is reversed, has proven to be a simple and effective technique for mitigating the position bias. Prior research (e.g., Loya et al., 2023) also found that LLMs are sensitive to prompt variations. Wei et al. (2022) demonstrated that chain-of-thought prompting resulted significantly enhanced LLMs performance on complex reasoning tasks. However, whether prompt variations similarly influence the performance of LLMs in paired comparisons remains an open empirical question.

In educational measurement, several studies have used the paired comparison method with LLMs as judges to estimate item difficulty. For example, Ballon et al. (2025) used OpenAI's o3 and Gemini 2.5 Pro to make paired comparisons on items that measure advanced math (34 items), English reading (787 items), and algebra (1876 items), respectively. They found the difficulty estimates based on the paired comparisons of Gemini-2.5-pro correlated with empirical difficulty at .44 for the advanced math items and .53 for English reading items, while OpenAI o3 generated weaker correlations. Although the correlations were not strong, this study demonstrated the promise of the paired comparison approach for estimating item difficulty. However, the study did not account for LLM position bias or provide sufficient details on how the empirical item difficulties were obtained.

Kolesnikova et al. (2026) examined the performance of three LLMs, GPT-4o, DeepSeek-Chat, and Qwen3-235b, in direct difficulty estimation and pairwise comparisons of 360 mathematics items measuring six domains [1] of elementary school arithmetic operations and applications. Results showed an average Spearman correlation of .73 across domains and LLMs, using paired comparisons with a zero-shot prompting and a hard decision rule (i.e., based on explicit outputs of the LLMs). They also found that pairwise judgments yielded better correlations than absolute judgments by LLMs, and that incorporating few-shot prompting or token-level probability failed to outperform zero-shot prompting or the hard decision rule. However, as with the study by Ballon et al. (2025), the study did not control for position

[1] Paired comparisons were made separately for each domain.

bias of the LLMs, so it was possible that the reported correlations were inflated.

Li and Gushta (2026a) selected five LLMs and prompted them to serve as judges and make pairwise comparisons of the difficulty of 40 Grade 5 English reading items. They examined within-model and between-model reliability of the ratings and the effects of model temperature and number of comparisons per item on rating accuracy and item difficulty recovery. Results showed that low temperature led to high between- and within-model reliability. Item difficulty estimates based on LLM judges' paired comparisons correlated moderately to strongly with empirical item difficulty, with the highest coefficient at .62 based on the ratings by Gemini-3.1-flash. This study provided many useful implications for practitioners, but it used a small set of reading items and a single prompt. It remains to be seen how LLMs would perform with items from a different content domain, a larger item set, and different prompts.

This study aims to examine the performance of LLM in making paired comparisons to estimate the difficulty of mathematics items developed for Grades 3-8 students. We selected three LLMs and designed several prompts to identify the best performing model and the most effective prompt. We also investigated factors that affect the rating accuracy of LLMs by using empirical item difficulty as the ground truth. Specifically, we intend to answer the following research questions:

1. How well are mathematics item difficulties estimated using paired comparisons of LLMs?
2. Does prompt design affect LLMs' performance in making paired comparisons?
3. Are there any item-pair related features that affect the rating accuracy of the LLMs in paired comparisons?

3 Methods

3.1 Mathematics Items

A total of 178 mathematics items developed for students in Grades 3-8 in the U.S. were used for the study. These items were from an empirical study in which 11,786 students in Grades 3-8 took a fixed-form mathematics test. Test lengths ranged from 30 to 36 items, depending on grade levels. Common items were designed for vertical linking, resulting in 178 unique items in the pool. These items were intended to measure the following content domains: (1) Number Sense; (2) Numerical Operations; (3) Algebra and Algebraic Thinking; (4) Geometry; (5) Measurement; and (6) Data Analysis, Statistics, and Probability (MetaMetrics, 2022). All items are multiple-choice, each consisting of a stem and four options. The empirical difficulty estimates of the items were obtained by applying the Rasch model (Rasch, 1960/1980) to students' item response data pooled across grades. Consequently, the estimated difficulties of all items are on the same scale. A sample item is shown below:

Felipe purchased a suit on sale for $131.25. The sale price was 30% off of the original price. Which was the original price of the suit?

A. $91.88
B. $170.63
C. $187.50 (correct option)
D. $437.50

The sample item above only contains text and numbers. However, most items have mathematical symbols, and about 39% of the 178 items have images in the stem, the options, or both.

3.2 Large Language Models

Although a wide range of LLMs are available, due to the characteristics of the items used in this study, we applied the following criteria when selecting LLMs: (1) the model is capable of processing both text and image data; (2) the model has strong mathematical reasoning capability; and (3) the model is fast and cost-effective. Using these criteria, we selected three LLMs: Gemini-3.5-flash, Claude-sonnet-4.6, and GPT-5.4.

3.3 Experimental Design

For Gemini-3.5-flash and Claude-sonnet-4.6, the temperature was set to 0 to minimize randomness in their responses and maximize within-model rating reliability. GPT-5.4 does not allow temperature setting so the default 1 was used. To control the position bias of LLMs, we used the swapping method (Zheng et al., 2023) to switch the positions of the items in each pair and instructed the LLM to compare the same pair in both orders. If the selected items from the two comparisons for the same item pair were the same, it was recorded as the final rating for that pair; otherwise, it was a tie. The tie also reflects models' position bias.

With N items, the total number of unique item pairs is N*(N-1)/2. Based on the findings by Li and

Gushta (2026a) and following the guidelines by Jones and Davies (2023), 10 comparisons per item were executed. Only 33 Grade 5 items were used to study the prompt sensitivity of LLMs and identify the best-performing model and the most effective prompt. Grade 5 items were selected because they are considered most representative of the entire Grades 3-8 item pool.

All models had the same system prompt or persona (see the appendix). To examine prompt sensitivity of LLMs, we designed three user prompts: (1) a prompt that simply instructed the LLM to compare the difficulty of two items in each pair and choose the one that is more difficult for students; (2) a prompt that includes everything in Prompt 1 but added vague evaluation criteria (i.e., criterion with names but no descriptions) for the model to consider when comparing item difficulty; and (3) a prompt containing detailed evaluation criteria, each accompanied by a description and illustrative examples (see the appendix). For each item, the stem, options, and key were provided as input to the LLM. For items with images, the images files (save in PNG format) were included as part of the input. Each LLM served as a judge and was instructed to complete all paired comparisons for each prompt. All models were configured to operate at a "high" reasoning or thinking level.

Paired comparison data using Grade 5 items were collected between June 11 and June 14, 2026. The selected model and prompt were then used on all 178 items. The paired comparison data for all items was collected between August 28 and August 30, 2026.

3.4 Data Analysis

After data collection, the paired comparison data was analyzed by applying the Bradley-Terry model (Bradley & Terry, 1952) using Facets (Linacre, 2023a). Using the model, the probability of an item i being judged as more difficult than item j is:

$$Prob\ (i > j) = \frac{e^{\beta_i}}{e^{\beta_i + e^{\beta_j}}}, \qquad (1)$$

where β_i and β_j are item difficulty parameters on the logit scale. Joint maximum likelihood method was implemented in Facets to estimate item parameters, assuming local independence among pairwise comparisons.

After obtaining initial item difficulty estimates from Facets, we observed that they had a large spread, which is often caused by data sparseness, a Guttman-like ordering of the items, or the presence of extremely easy or difficult items (Linacre, 2023b). For a fair comparison between prompts and LLMs, we applied linear transformation to item difficulty estimates based LLM judges' paired comparisons such that the transformed estimates had the same mean and standard deviation as the empirical difficulty estimates.

To evaluate item difficulty recovery, we computed the Pearson correlation between LLM-paired-comparison-based difficulty estimates and empirical difficulty estimates of 33 Grade 5 items and the root mean squared error (RMSE) of LLM-derived difficulty. The correlations and RMSEs were compared across prompts and models to determine the most effective prompt and best-performing model. The prompt or the model yielding the strongest correlation and the smallest RMSE was designated as optimal. The best-performing LLM with the most effective prompt was then used with the entire 178 items in the Grades 3-8 item pool.

To evaluate prompt sensitivity of the LLMs, we compared item difficulty recovery across prompts and then performed a chi-square independence test on the rating accuracy of each model with different prompts. Rating accuracy was computed by dividing the number of times the correct item was selected in a pair (judged by their empirical difficulty) by the total number of item pairs compared. When a tie occurred, the accuracy rate for the item was halved. Additionally, we also computed the tie rate to record position bias of the models. The analyses were performed for Grade 5 items only.

To answer the third research question, we fit a logistic regression model to the paired comparisons of all 178 items using rating accuracy as the outcome variable and three identified predictors. The rating accuracy had two values, with 0 indicating the wrong item was selected and 1 otherwise. The three predictors are: (1) whether the two items were developed for the same grade level; (2) whether at least one item in the pair has images in its stem or options, and (3) whether the two items measure the same content strand. We hypothesized that the first predictor would contribute most to rating accuracy, given that difficulty of items is typically strongly correlated with their associated grade level. The effect of image on the rating accuracy was examined because preliminary testing revealed that certain LLMs produced incorrect answers to items with images in the stem

or options during exploration. Difference in content domains could be another factor affecting accuracy because a previous study (Li & Gushta, 2026b) found that construct heterogeneity can influence human judges' paired comparisons of reading items.

4 Results

Table 1 shows item difficulty recovery results using 33 Grade 5 items. Among the three LLMs, Gemini-3.5-flash yielded the highest correlations between LLM-derived item difficulty and their empirical counterparts, with coefficients around .70 across prompts. GPT-5.4 and Claude-4.6-sonnet generated lower correlations between the two sets of item difficulties, with coefficients around .60. The 95% confidence intervals of the correlation coefficients are wide, reflecting the small number of items in the analyses. With regard to RMSE, Gemini-3.5-flash was also the winner because of its lowest RMSEs in the item difficulty estimate among the three LLMs. The other two LLMs had comparable RMSEs. Figure 1 shows the scatter plot between LLM-derived difficulty and empirical difficulty for 178 Grades 3-8 items based on the paired comparisons of Gemini-3.5-flash with a simple prompt. A strong, positive, and linear relationship was observed between the two sets of item difficulties. The correlation coefficient was .80 with 95% confidence interval between .74 and .85. The RMSE of the item difficulty estimates was .77.

In terms of prompt sensitivity of the LLMs, there appears to be a weak interaction effect between

LLM	Prompt	*r* [95% CI]	RMSE
Claude-4.6-sonnet	Simple	.62 [.35, .79]	.87
	Vague	.57 [.28, .76]	.93
	Detailed	.56 [.27, .76]	.93
Gemini-3.5-flash	Simple	.70 [.48, .84]	.77
	Vague	.69 [.46, .84]	.78
	Detailed	.71 [.48, .85]	.76
GPT-5.4	Simple	.59 [.31, .77]	.91
	Vague	.62 [.36, .80]	.86
	Detailed	.56 [.26, .76]	.94

Table 1: Item difficulty recovery of Grade 5 items.

prompt type and LLM: Gemini-3.5-flash appeared insensitive to different prompts, but a simple prompt led to marginally better performance of Claude-4.6-sonnet, while a prompt with vague evaluation criteria for item difficulty comparison

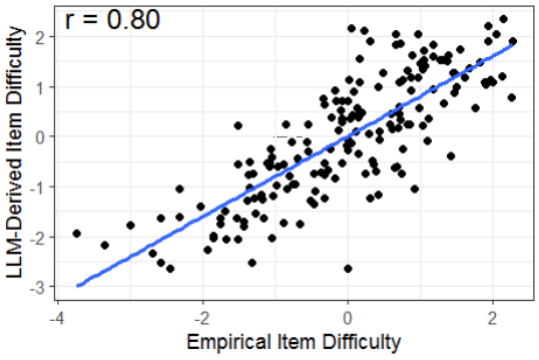


Figure 1: Scatter plot of LLM-derived and empirical difficulty of 178 Grades 3-8 mathematics items.

resulted in slightly better performance by GPT-5.4.

Table 2 shows the position bias (indicated by tie rate) and rating accuracy of the LLMs and the chis-square statistic for rating accuracy. Gemini-3.5-flash showed the least amount of position bias among the three LLMs, followed closely by GPT-5.4, while Claude-4.6-Sonnet had the highest position bias. Similarly, Gemini-3.5-flash also had the highest accuracy rate in its ratings, followed by GPT-5.4 and Claude-4.6-Sonnet. The chi-square independence tests were not significant for any prompts, suggesting that prompt type had little effect on the rating accuracy of the LLMs.

LLM	Prompt	Tie %	Accuracy %	χ^2
Claude-4.6-sonnet	Simple	24	56	.76
	Vague	22	55	
	Detailed	18	58	
Gemini-3.5-flash	Simple	9	66	.04
	Vague	11	66	
	Detailed	11	67	
GPT-5.4	Simple	11	62	.14
	Vague	13	63	
	Detailed	10	63	

Table 2: Position bias (tie rate) and rating accuracy of LLMs across prompts based on 33 Grade 5 items.

The effects of the three selected factors on the rating accuracy of Gemini3.5-flash with a simple prompt and 178 mathematics items are displayed in Table 3. Two factors, whether two items have the same assigned grade levels and whether at least one item has images in a pair had statistically significant effects on the rating accuracy. *grade_same* had a significant negative effect, indicating that same grade levels resulted lower

rating accuracy than different grade levels. When at least one item has images in its stem, options, or both, it had a positive effect on the rating accuracy. Whether two items in a pair measure the same content strand had little effect on the rating accuracy, as indicated by the nonsignificant coefficient (p = .664) and the odds ratio close to 1.

Variable	$\hat{\beta}$	SE	exp($\hat{\beta}$)	Z
Intercept	1.08	0.11	2.93	9.48***
grade_same	-0.55	0.15	0.58	-3.65***
with_image	0.41	0.13	1.50	3.24**
strand_same	0.07	0.15	1.07	0.44

Note. *** p < .001; ** p < .01.

Table 3: Results of logistic regression analysis of rating accuracy of Gemini-3.5-flash.

5 Discussion

This study investigates the effectiveness of using LLM judges' paired comparisons to estimate mathematics item difficulty. The results for the 33 Grade 5 mathematics items showed that the correlations between LLM-paired-comparison-based difficulty estimates and empirical difficulty estimates were moderate to strong, with Gemini-flahs-3.5 achieving the best performance in item difficulty recovery among the three LLMs. Ballon et al. (2025) also found that Gemini-2.5-Pro performed better than OpenAI's o3. The better performance of Gemini might be attributed to its stronger capability of processing image data (Kokalj, 2026). When using the paired comparisons of Gemini-flash-3.5 on the entire 178 Grades 3-8 mathematics items with a simple prompt, the correlation coefficient reached .80. The correlation was quite strong, only slightly weaker than the highest correlation reported in Razavi and Powers' (2025) study, but that study used 2,564 mathematics items and considerable effort and resources were invested in extracting item features. Also, it is also worth noting that only 10 comparisons per item were used for the 178 items.

Regarding the prompt sensitivity of the LLMs, the patterns were slightly different among the models using item difficulty correlations and RMSE as the criteria. For Gemini-3.5-flash, minimal difference in correlation or RMSE were observed among prompts. For Claude-4.6-sonnet, the simple prompt resulted in better item difficulty recovery, while for GPT-5.4, a prompt with vague evaluation criteria demonstrated a marginal advantage. We did not find statistically significant differences among the three prompts in the rating accuracy across LLMs. Overall, no strong evidence was found to support that providing detailed comparison criteria in the prompt increases the performance of LLMs in paired comparison of mathematics items. The results were probably due to the fact that the reasoning levels of all the large language models used were already high. That is, even without providing any evaluation criteria, the LLMs already used a set of criteria similar to the criteria provided in the other two prompts to compare item difficulties.

In terms of item features that contribute to the rating accuracy of the LLMs, two significant predictors were identified: whether the intended item grade levels are the same and whether at least one item in a pair contains images. Given the strong correlations observed between LLM-based and empirical difficulty, it was expected to see whether item grade levels are the same in a pair was a significant predictor. Items with the same grade level tend to be more similar in difficulty than items with different grade levels. The higher similarity in difficulty for items with the same grade level should have made it harder for Gemini-3.5-flash to determine the correct item in a pair than items with larger differences in difficulty. The second factor, whether at least one item has an image in its stem or options, had a positive effect on the rating accuracy, suggesting that when this condition is true, Gemini-3.5-flash has better performance in paired comparison of mathematics item difficulty. Whether the content strands are the same for items in a pair had little effect on the rating accuracy. This was likely because the measured strands have a clear progression in curriculum and instruction; mathematics skills and knowledge taught later to students inherently build on foundational concepts introduced earlier. Therefore, whether two items in a pair measure the same content strands had little impact on the rating accuracy of the model.

The findings of the study has important implications for practitioners. It shows that using paired comparison with LLMs can generate difficulty estimates of mathematics items that strongly correlated with empirical item difficulties. These correlations are strong enough to be useful for low-stakes classroom assessment, personalized learning, and adaptive testing. The correlations are higher than many of the correlations found using alternative human-response-data-free approaches

to item difficulty estimation and comparable to those using resource-intensive data-driven approaches. Additionally, compared to alternative approaches, it is more efficient, cost-effective, and easier to implement. Furthermore, using LLMs as judges with the paired comparison method does not require items with existing difficulty. This is especially valuable for small or new testing programs where it is not feasible to conduct field testing to obtain initial estimates of item difficulty.

This study has several limitations. First, in order to use LLMs that were roughly comparable in performance, we chose the three models. However, due to the unique features of the models, their temperatures settings differed, making it difficult to ensure a fair comparison across models. Second, our model selection was constrained by cost efficiency. Future research with access to greater computation resources should evaluate larger or more advanced models, which may yield more accurate item difficulty estimates. Third, we only implemented 10 comparisons per item. Li and Gushta's (2026a) study found little difference in item difficulty recovery between 10 and 15 comparisons per item. However, that study was based on a small set of items. It is possible that more comparisons per item would result in better item difficulty estimation for a larger item set. This could be an experiment design factor for future studies. Fourth, we did not examine between- or within-model reliability of the ratings because we assumed that the findings from Li and Gushta (2026a) study based on reading items generalize to mathematics items. Given the critical role of reliability in paired comparison, evaluating this assumption is particularly important when the model temperature is high. Fifth, we did not prompt LLMs to generate justifications or rationale for their choices. Eliciting qualitative explanations could yield valuable insights into LLM reasoning during pairwise evaluations, ultimately informing strategies to optimize prediction accuracy.

References

AlKhuzaey, S., Grasso, F., Payne, T. R., & Tamma, V. (2024). Text-based question difficulty prediction: A systematic review of automatic approaches. *International Journal of Artificial Intelligence in Education*, *34*(3), 862-914. https://doi.org/10.1007/s40593-023-00362-1

Attali, Y., Saldivia, L., Jackson, C., Schuppan, F., & Wanamaker, W. (2014). *Estimating item difficulty with comparative judgments* (ETS Research Report No. RR-14-39). Princeton, NJ: Educational Testing Service. https://files.eric.ed.gov/fulltext/EJ1109277.pdf

Ballon, M., Algaba, A., Verbeken, B., & Ginis, V. (2025). *Estimating problem difficulty without ground truth using large language model comparisons*. arXiv. https://doi.org/10.48550/arXiv.2512.14220

Bejar, I. I. (1983). Subject matter experts' assessment of item statistics. *Applied Psychological Measurement*, *7*(3), 303-310. https://doi.org/10.1002/j.2333-8504.1981.tb01274.x

Bradley, R. A., & Terry, M. E. (1952). Rank analysis of incomplete block designs: I. The method of paired comparisons. *Biometrika*, *39*(3/4), 324. https://doi.org/10.2307/2334029

Cai, H., Yan, X., Chuang, P. L., Pan, Y., & Huo, M. (2026). What makes listening comprehension difficult?: A feature-based machine learning approach to understanding item difficulty. *Applied Linguistics*, *47*(2), 406-428. https://doi.org/10.1093/applin/amaf079

He-Yueya, Joy, Wanjing Anya Ma, Kanishk Gandhi, Benjamin W. Domingue, Emma Brunskill, and Noah D. Goodman. (2024). *Psychometric alignment: Capturing human knowledge distributions via language models*. https://arxiv.org/abs/2407.15645

Huang, Z., Liu, Q., Chen, E., Zhao, H., Gao, M., Wei, S., Su, Y., & Hu, G. (2017). Question difficulty prediction for reading problems in standard tests. In *Proceedings of the AAAI Conference on Artificial Intelligence* (pp. 1352-1359). AAAI Press. https://doi.org/10.1609/aaai.v31i1.10740.

Jones, I., & Davies, B. (2023). Comparative judgement in education research. *International Journal of Research & Method in Education*, *47*(2), 170-181. https://doi.org/10.1080/1743727X.2023.2242273

Kokalj, E. (2026, May 22). *Gemini 3.5 Flash for vision: Evaluation and benchmarks.* Roboflow. https://blog.roboflow.com/use-gemini-3-5-flash-vision/?utm_source=chatgpt.com

Kolesnikova, D., Fedyanin, K., Hofman, A. D., Brinkhuis, M. J. S., & Bolsinova, M. (2026). *Estimating item difficulty with large language models as experts*. arXiv. https://doi.org/10.48550/arXiv.2605.18562

Li, L., & Gushta, M. (2026a, April 8-11). *Using AI judges' paired comparisons to estimate item difficulty* [Paper presentation]. National Council on Measurement in Education (NCME) Annual Meeting, Los Angeles, CA, United States.

Li, L., & Gushta, M. (2026b). *Using content experts' paired comparison to estimate item Difficulty: What affects the accuracy?* [Manuscript under review]. Department of Research and Development, MetaMetrics.

Linacre, J. M. (2023a). *Facets computer program for many-facet Rasch measurement* (Version 3.86.0) [Computer software]. Winsteps.com.

Linacre, J. M. (2023b). *Facets (64-bit) help for Many-Facets Rasch measurement.* https://www.winsteps.com/facetman64/index.htm?pairedcomparisons.htm

Liu, Y., Zhou, H., Guo, Z., Shareghi, E., Vulić, I., Korhonen, A., & Collier, N. (2024). *Aligning with human judgement: The role of pairwise preference in large language model evaluators*. arXiv. http://arxiv.org/abs/2403.16950

Liusie, A., Manakul, P., & Gales, M. (2024). LLM comparative assessment: Zero-shot NLG evaluation through pairwise comparisons using large language models. In Y. Graham & M. Purver (Eds.), *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)* (pp. 139–151). Association for Computational Linguistics. https://doi.org/10.18653/v1/2024.eacl-long.8

Lorge, I., & Kruglov, L. (1952). A suggested technique for the improvement of difficulty prediction of test items. *Educational and Psychological Measurement, 12*(4), 554-561. https://doi.org/10.1177/001316445201200402

Loya, M., Sinha, D., & Futrell, R. (2023). Exploring the sensitivity of LLMs' decision-making capabilities: Insights from prompt variations and hyperparameters. In H. Bouamor, J. Pino, & K. Bali (Eds.), *Findings of the Association for Computational Linguistics: EMNLP 2023* (pp. 3711–3716). Association for Computational Linguistics. https://doi.org/10.18653/v1/2023.findings-emnlp.241

Maeda, H. (2025). Field-testing multiple-choice questions with AI examinees: English grammar items. *Educational and Psychological Measurement, 85*(2), 221-244. https://doi.org/10.1177/00131644241281053

MetaMetrics. (2022). *Quantile Framework for Mathematics: Development and validity evidence.* https://partnerhelp.metametricsinc.com/image/Quantile%20Framework%20for%20Mathematics%20Validity%20Evidence_2022.pdf

Qin, Z., Jagerman, R., Hui, K., Zhuang, H., Wu, J., Yan, L., Shen, J., Liu, T., Liu, J., Metzler, D., Wang, X., & Bendersky, M. (2023). *Large language models are effective text rankers with pairwise ranking prompting*. arXiv. https://arxiv.org/abs/2306.17563

Rasch, G. (1960). *Probabilistic models for some intelligence and attainment tests*, Copenhagen: Danish Institute for Educational Research (reprint 1980, expanded edition, with foreword and afterword by B. D. Wright). Chicago: The University of Chicago Press.

Razavi, P., & Powers, S. (2025). *Estimating item difficulty using large language models and tree-based machine learning algorithms*. arXiv. https://arxiv.org/abs/2504.08804

Rupp, A. A., Garcia, P., & Jamieson, J. (2001). Combining multiple regression and CART to understand difficulty in second language reading and listening comprehension test items. *International Journal of Testing*, *1*(3–4), 185–216. https://doi.org/10.1080/15305058.2001.9669470

Thurstone, L. L. (1927). The method of paired comparisons for social values. *The Journal of Abnormal and Social Psychology*, *21*(4), 384–400. https://doi.org/10.1037/h0065439

Wang, P., Li, L., Chen, L., Cai, Z., Zhu, D., Lin, B., Cao, Y., Liu, Q., Liu, T., & Sui, Z. (2023). *Large language models are not fair evaluators*. arXiv.Org. https://arxiv.org/abs/2305.17926

Wei, J., Wang, X., Schuurmans, D., Bosma, M., Ichter, B., Xia, F., Chi, E. H., Le, Q. V., & Zhou, D. (2022). Chain-of-thought prompting elicits reasoning in large language models. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, & A. Oh (Eds.), *Advances in Neural Information Processing Systems* (Vol. 35, pp. 24824–24837). Curran Associates, Inc. https://proceedings.neurips.cc/paper_files/paper/2022/hash/9d5609613524ecf4f15af0f7b31abca4-Abstract-Conference.html

Zheng, L., Chiang, W.-L., Sheng, Y., Zhuang, S., Wu, Z., Zhuang, Y., Lin, Z., Li, Z., Li, D., Xing, E. P., Zhang, H., Gonzalez, J. E., & Stoica, I. (2023). Judging LLM-as-a-judge with MT-bench and Chatbot Arena. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, & S. Levine (Eds.), *Advances in Neural Information Processing Systems* (Vol. 36, pp. 46595–46623). Curran Associates, Inc. https://proceedings.neurips.cc/paper_files/paper/2023/hash/91f18a1287b3e90d03fdf04da1011882-Abstract-Conference.html

A Appendix: LLM Prompts

System prompt (Persona): "You are a content expert in K-12 mathematics, specializing in item development and difficulty evaluation for math assessments."

Simple user prompt: "You are a content expert in K-12 mathematics and a math teacher in the United States. You specialize in developing math test items. You will be provided two mathematics items developed for Grades 5 students. They may contain text, images, or both. Compare the two items and determine which one is more difficult for students. Provide your answer using one word, that is, either 'item1' or 'item2'."

User prompt with vague criteria: "You will be shown two multiple-choice mathematics items developed for Grade 5 students. Each item contains a stem and four options, which may have or be images. Compare the two items and determine which one is more difficult for students. Consider cognitive demand and reasoning depth, linguistic and textual features, symbolic complexity, and distractor quality of the items. Before answering, think through the comparison step by step using the criteria above. Do this reasoning internally — do not include it in your response. Provide your response using one word, that is, either 'item1' or 'item2'."

User prompt with detailed criteria: "You will be shown two multiple-choice mathematics items developed for Grade 5 students. Each item contains a stem and four options, which may have or be images. Compare the two items and determine which one is more difficult for Grade 5 students. Consider the following criteria when comparing the two items.

Cognitive demand:

(1) Items requiring a sequence of operations or distinct lines of reasoning consistently increase difficulty;

(2) Problems that demand conceptual understanding, interpreting unstated information, or mathematical modeling are harder than rote calculation.

(3) Contextualized problems (word problems) add cognitive load compared to decontextualized, purely numerical equations

Linguistic and textual features:

(1) Word length and syntactic complexity are major predictors. High vocabulary complexity or long sentences force students to translate a reading task into a math task.

(2) Extraneous information (distractor sentences or extraneous numbers within a word problem) makes it harder for students to identify the correct algorithm.

Symbolic Complexity and Format:

(1) Items heavily laden with complex formulas, fractions, or algebraic notation are inherently more difficult than those using whole numbers.

(2) In early and intermediate math, certain arithmetic operations act as heavy predictors. Concepts such as borrowing, carry-over, or inversed ordering naturally increase cognitive burden

(3) The presence of graphs, diagrams, or geometry can increase or decrease difficulty. If the visual is critical to solving the problem, it adds spatial reasoning demand.

Distractor Quality:

Plausible but incorrect distractors drastically drive up difficulty. Conversely, easily eliminated or poorly crafted distractors (like implausible numbers or giveaways) make multiple-choice questions artificially easy.

Before answering, think through the comparison step by step using the criteria above. Do this reasoning internally—do not include it in your response. Provide your response using one word, that is, either 'item1' or 'item2'."

Math-Specialized or General-Purpose? A Comparison of MathBERT and DeBERTa-v3-large for Automated Scoring of Mathematics-Explanation Items

Gregory M. Jacobs **Ahmed H. Bediwy** **Martha Bellows**
Pearson Education
{greg.jacobs, ahmed.bediwy, martha.bellows}@pearson.com

Abstract

Automated scoring of math-explanation items should handle responses that mix natural language with symbolic reasoning. The 2023 NAEP Automated Scoring Challenge showed that fine-tuning pre-trained encoders can reach human-like agreement, but its winning system used a large, general-purpose encoder, while smaller math-specific pre-trained encoders remain a plausible and more efficient alternative. We fine-tune MathBERT (110M parameters) and DeBERTa-v3-large (183M parameters) as single-output regression scorers on 31 math-explanation items from one statewide summative assessment administration, benchmarking every model against human–human reliability and summarizing deployability with a three-tier acceptance criteria. On the 18 items fine-tuned under both encoders with identical datasets, the resulting performance between the two was practically indistinguishable: mean best QWK was 0.931 for MathBERT and 0.933 for DeBERTa-v3-large—a difference of only 0.002—and the models traded top performance at similar rates. Given similar performance, we tested a two-stage approach that used the larger general-purpose encoder only when the smaller math-focused encoder failed acceptance criteria. Extending MathBERT to all 31 items, it reached a mean QWK of 0.926 against a human–human benchmark of QWK = 0.940 and met operational acceptance criteria on 81% of items. Escalating the six items where a fine-tuned MathBERT model failed acceptance to a fine-tuned DeBERTa-v3-large model recovered four of the six, yielding acceptable automated-scoring models for 29 of 31 items. Because the smaller, math-pre-trained encoder matches the larger one at roughly 40% fewer parameters, we recommend fine-tuning a MathBERT encoder as the default and escalating to the larger, general purpose encoder like DeBERTa-v3-large only when a MathBERT model fails when deploying automated scoring models operationally at scale for math-explanation items.

1 Introduction

Constructed-response items that ask students to *explain* their mathematical reasoning yield valuable evidence about conceptual understanding, but hand-scoring them at the scale of a statewide summative assessment is resource intensive and time consuming. Automated scoring seeks to reduce those costs and streamline the scoring timeline. Math explanation items, however, are challenging to score automatically (Hellman et al., 2023; Morris et al., 2025b). A single response can interweave natural-language argumentation, symbolic notation, procedural steps, and domain-specific vocabulary, so a model must reason over a mixture of forms that general-purpose encoders were not designed to represent. The difficulty is compounded by what these items ask for. Students must supply both a mathematically correct solution and an account of the reasoning that produced it, so a scoring model must evaluate both at once (Hellman et al., 2023).

The 2023 NAEP Mathematics Automated Scoring Challenge established an important proof of concept: fine-tuned transformer models can reach *human-like* agreement—model–human quadratic weighted kappa within roughly 0.05 of human–human reliability—on most math-explanation items (Crossley et al., 2023). The winning system was built on DeBERTa-v3-large, a 183M-parameter general-purpose encoder (He et al., 2021, 2023), while the runner-up used a math-specialized RoBERTa variant (NAEP Research and Development Program, 2023). The broader literature highlights that smaller, math-pretrained encoders such as MathBERT (Shen et al., 2021) are potential efficient alternatives. This raises a question the challenge did not resolve: should a smaller math-specialized encoder, rather than a larger general-purpose one, be used to score math-explanation items operationally at scale?

This paper seeks to answer that question. We

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 193–199
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

fine-tune two encoders, a smaller specialized MathBERT encoder (110M parameters) and a larger general-purpose DeBERTa-v3-large encoder (183M), as single-output regression scorers. To keep the comparison interpretable, we hold the assessment context fixed. The corpus is 31 math-explanation items from one administration of a single statewide summative assessment.

We contribute a head-to-head comparison of a math-specialized and a general-purpose transformer on operational data, under a protocol that references every model against human–human reliability and summarizes deployability with a three-tier agreement-to-standard flag that programs can adopt directly. Concretely, we ask: (RQ1) Does a fine-tuned math-specialized encoder match the performance of a larger general-purpose encoder on math-explanation items? and (RQ2) Does escalating to a larger general-purpose encoder for items where a fine-tuned math-specialized encoder does not meeting acceptance criteria improve performance? Based on the answer to these questions, we propose a training pipeline that can be deployed to fine-tune encoders for automated scoring of operational math explanation items at scale.

2 Related Work

Automated scoring of constructed responses: Established evaluation frameworks in educational measurement hold that an engine's usefulness should be judged not by raw accuracy but by its agreement with human raters *relative to* the agreement humans achieve with one another, together with fairness and score-distribution considerations (Bennett and Zhang, 2011; Pearson and ETS, 2015; Williamson et al., 2012; Yan et al., 2020). Two properties of operational data shape our method. First, rubric scores are ordinal, so regression under a mean-squared-error objective can match or exceed cross-entropy classification because it respects the ordering of score points (Morris et al., 2025a). Second, severe class imbalance is endemic (Crossley et al., 2023; Yan et al., 2020).

Transformer encoders and domain-adaptive pretraining: The transformer architecture (Vaswani et al., 2017) and its bidirectional encoder instantiation BERT (Devlin et al., 2019) underpin most modern scoring systems. DeBERTa introduced disentangled attention (He et al., 2021), and DeBERTa-v3 improved pre-training with an ELECTRA-style objective (He et al., 2023). In parallel, encoders pretrained on mathematical corpora—MathBERT (Shen et al., 2021) and math-specialized RoBERTa variants (NAEP Research and Development Program, 2023)—aim to capture the notation and discourse of mathematics with far fewer parameters. Our comparison is a controlled test of the tradeoff: MathBERT embodies the domain-pretraining approach and DeBERTa-v3-large the scale approach, fine-tuned under an identical harness so that only the encoder choice differs.

The NAEP challenge and operational deployment: The top-performing systems in the NAEP challenge were optimized for peak accuracy on a fixed benchmark, using extensive data-augmentation and ensembling pipelines (Crossley et al., 2023; NAEP Research and Development Program, 2023). Operational programs face a different objective: sustaining human-like agreement at acceptable performance metrics across many items and administrations in an efficient manner, with score distributions close enough to those of human raters that reported results stay comparable year over year. We focus explicitly on identifying the optimal training pipeline for operational programs.

3 Corpus

3.1 Data

Every statistic in this paper is computed from a single administration of 31 math-explanation items from one statewide summative assessment. The items span from grades 3 through 8 and apply rubrics with scorepoint ranges that span from two (0-1) up to five (0-4) integer values.

3.2 Human scoring

Responses were scored under an operational double-read procedure whereby 10% of responses received a second score to measure human-to-human reliability. For both training and evaluation, we take the target label to be `best_human`: the program's final reported score, determined by a predefined hierarchy in which an adjudication or resolution score takes precedence over the first and second reads. Every metric we report is computed on a held-out validation set of 2,500 responses per item. Human–human reliability is taken from the first and second reads on those same set of responses, so the human metrics are calculated on the smaller 10% of those responses that received a second human score (approximately 250 validation responses per item).

3.3 Response pre-processing

Operational responses arrive in a semi-structured delivery format rather than as clean text, and they carry markup that a language model should not see verbatim. Before fine-tuning, we apply a deterministic cleaning pipeline: we unwrap the JSON delivery format, decode HTML entities and strip presentational tags, convert embedded MathML into a plain-text rendering so that symbols, fractions, and operators survive as readable tokens, and normalize Unicode. The pipeline is identical across all items and both encoders, so no difference between models can be attributed to pre-processing.

4 Modeling Methodology

4.1 Data splits

For each item we drew a random sample of 7,500 responses and stratified it into 5,000 for training—five folds of 1,000 for cross-validated hyperparameter selection—and 2,500 for validation. Stratification preserves the score-point distribution across splits, and all models for an item share the same splits, so comparisons are paired.

4.2 Transformers implementation

Both encoders use `transformers` (Wolf et al., 2020) with each model's native tokenizer, AdamW (Loshchilov and Hutter, 2019) under a linear warmup schedule, a physical batch size of 4 with four gradient accumulations (effective batch size of 16), and early stopping on validation QWK. The harness is shared, so architectural rather than harness differences drive any performance gap.

4.3 Hyperparameter grid search

Hyperparameter values are identified using the training dataset to perform a five-fold stratified cross-validation grid search, varying learning rate, weight decay, and label-smoothing α parameter values and evaluating each by mean validation QWK. Once the optimal hyperparameter values are identified through the grid search, a final model is trained using those optimal parameter values.

4.4 Model fine-tuning

Each encoder is fine-tuned with a single-output regression head, which respects the ordinal structure of rubric scores (Morris et al., 2025a). To counter imbalance, we minimize a sample-weighted mean-squared-error loss with class-balanced label smoothing. Let $\hat{y}_i$ be the model's continuous prediction for response i, let y_i be that response's human score on a rubric with score points $\{0, \ldots, K-1\}$, and let f_k be the empirical frequency of score point k in the training set. With per-response weight w_i and smoothing coefficient $\alpha \in [0, 1]$,

$$
\begin{aligned}
w_i &= (1-\alpha)\,\frac{1/f_{y_i}}{\frac{1}{K}\sum_k 1/f_k} + \alpha, \\
\mathcal{L} &= \frac{1}{N}\sum_{i=1}^{N} w_i\,(\hat{y}_i - y_i)^2,
\end{aligned}
\tag{1}
$$

which reduces to inverse-frequency (class-balanced) weighting at $\alpha = 1$ and to uniform weighting at $\alpha = 0$, up-weighting rare high-score responses without destabilizing training on the dominant low-score class. Smoothing α is treated as a hyperparameter because the right amount of correction varies with how skewed a item's score distribution is.

4.5 Evaluation and deployability

Models are evaluated with four complementary statistics, each benchmarked against the corresponding human–human value so that "good" always means "close to what two humans achieve." *Quadratic weighted kappa* is the primary statistic (Cohen, 1968): disagreements are weighted by squared distance, which suits ordinal scores. *Exact agreement* is the proportion of responses assigned the same score point by model and human. The *standardized mean difference* captures systematic leniency or severity. Finally, *score-point distribution match* compares model and human marginal frequencies, guarding against a model that agrees well while distorting the reported distribution and breaking year-over-year comparability.

These four statistics are combined into a three-tier deployability flag (Figure 1). A *green* flag requires that (i) model-minus-human QWK ≥ -0.10, (ii) the exact-agreement gap ≥ -0.0525, (iii) $|SMD| \leq 0.15$, and (iv) every score-point frequency falls within four percentage points of the human distribution. *Yellow* means the agreement metrics pass but a score-point frequency falls outside that band; *red* means at least one agreement metric is not met. The thresholds are *deltas* against each item's own human–human values, so an item on which humans agree strongly is held to a stricter standard.

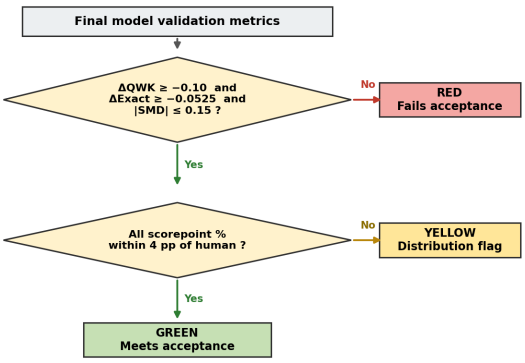


Figure 1: The three-tier agreement-to-standard flag. Green requires human-like agreement on all metrics and a close score-point distribution match.

These model evaluation metrics are considered industry standard and are consistently applied within the automated scoring field (Pearson and ETS, 2015; Williamson et al., 2012).

5 Initial Modeling Results

The initial round of modeling involved fine-tuning both MathBERT and DeBERTa-v3-large models, using a robust hyperparameter grid search, on only a subset of 18 of the 31 items to conserve resources. The purpose of this initial experimentation was to (1) evaluate the effectiveness of the robust hyperparameter grid search to determine whether a more limited search could be used at scale, and (2) compare the performance metrics between the two encoders to determine which one is more effective at scoring math-explanation items.

5.1 Hyperparameter grid search

For each encoder, we elected to search over learning rate, weight decay, and label-smoothing α and set the values for each based on the best performing fine-tuned models surveyed during the 2023 NAEP Mathematics Automated Scoring Challenge referenced above. Table 1 summarizes the hyperparameter values used in the grid search.

The main findings from this grid search results were that (1) the only hyperparameter that potentially should be tuned at scale for both encoders is learning rate, and (2) that the number of epochs should be increased with early stopping to prevent wasted resources. The "best" performing models for both encoders used the largest learning rate value a majority of the time (albeit with a marginally better QWK value), suggesting that a grid search with larger learning rate values could be implemented out of an abundance of caution.

Hyperparameter	Values searched
Learning rate	1e-5, 2e-5 (MathBERT); 1e-5, 1.5e-5, 2e-5 (DeBERTa)
Weight decay	0.01, 0.05 (MathBERT); 0.1, 0.3 (DeBERTa)
Smoothing α	0.5, 0.7
Batch size	4 (with 4 grad. accumulation)
Epochs	2
Scheduler	linear with warmup

Table 1: Stage 1 search space on 18 items. Grids yielded 8 (MathBERT) and 12 (DeBERTa-v3-large) configurations over 5 folds.

5.2 Head-to-head comparison

Table 2 compares the encoders on the 18 items where both models were run using identical splits, scoring each by its *best* configuration. Mean best QWK was 0.931 for MathBERT and 0.933 for DeBERTa-v3-large—a negligible difference of 0.002. The models traded wins at similar rates, 8 for MathBERT to 10 for DeBERTa-v3-large, with a mean absolute per-item exact agreement gap of only 0.004. The model results on SMD was similarly equivalent with DeBERTa-v3-large winning by 0.04. The fine-tuned MathBERT models did generate a lower mean SMD value and passed the acceptance criteria at a marginally larger rate, suggesting that it can be fine-tuned to score math-explanation items with the same, if not greater, success rate as the larger DeBERTa-v3-large model.

6 Final Training Results

6.1 Proposed Operational Training Pipeline

Based on the initial training results, we executed an operational training pipeline across all 31 items using a revised, more targeted hyperparameter grid search focused on the learning rate hyperparameter and the number of training epochs (with early stopping implemented). Because in our initial experimentation both encoders performed slightly better at higher learning rate values, we chose to expand the grid search at the upper bounds. For MathBERT, we used learning rate values of 2e-5, 3e-5, 4e-5, and 5e-5, and allowed training to run for up to 6 epochs. For DeBERTa-v3-large, we used learning rate values of 1.5e-5, 2e-5, 3e-5, and 4e-5, and allowed training to run for up to 4 epochs.

Additionally, we implemented a two-stage training approach. For stage 1, we fine-tuned only Math-

Configuration	Items	Mean QWK	Mean exact	Mean $\lvert SMD \rvert$	% green
Human–human reference	31	0.940	0.902	—	—
Initial Modeling Results (18 items)					
MathBERT	18	0.931	0.852	0.036	78%
DeBERTa-v3-large	18	0.933	0.856	0.056	67%
Final Modeling Results (31 items)					
MathBERT	31	0.926	0.872	0.015	81%
DeBERTa-v3-large	6	0.917	0.836	0.013	67%

Table 2: Mean engine–human agreement relative to human–human reliability.

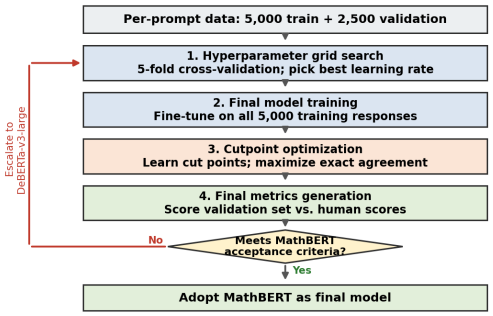


Figure 2: Per-item modeling pipeline. A model is adopted only if it meets the acceptance criteria; otherwise the program escalates to the larger encoder.

BERT models for all 31 items and evaluated performance. We selected MathBERT over DeBERTA-v3-large as our preferred encoder because it performed on par with the larger model in our initial experimentation, it can be fine-tuned more efficiently, and it is specifically designed to evaluate text that includes math notation. If a fine-tuned MathBERT model passed the acceptance criteria, we selected it as the model to be used operationally.

If a fine-tuned MathBERT model was unable to pass the acceptance criteria for an item, a second round of training was performed where we fine-tuned a DeBERTa-v3-large using the same training pipeline. If a fine-tuned DeBERTa-v3-large model passed the acceptance criteria, we selected it as the recommended model to be used operationally for the item. Figure 3 provides a visual representation of this workflow.

6.2 Final Modeling Performance

Fine-tuning a MathBERT model to meet acceptance criteria was successful for 25 of the 31 items, producing an 81% success rate. For the six items where a fine-tuned MathBERT model failed to pass acceptance criteria, all six failed because of the requirement by producing a model-to-human score agreement rate greater than 0.0525.

	Human–human	MathBERT	DeBERTa
Item 1	0.949	0.858	0.871
Item 2	0.822	0.762	0.738
Item 3	0.852	0.796	0.820
Item 4	0.874	0.798	0.863
Item 5	0.878	0.814	0.841
Item 6	0.909	0.850	0.884
Mean	0.881	0.813	0.836

Table 3: Exact agreement on the six items MathBERT failed to earn green, showing human–human reliability alongside MathBERT and the escalated DeBERTa-v3-large model agreement rates.

A second round of training to fine-tune DeBERTa-v3-large for these items produced a passing model for four of the six items. So, the proposed operational training pipeline was able to produce acceptable automated scoring models for 29 of 31 items in the subject administration. Table 4 outlines the exact agreement rates for humans and the best MathBERT and DeBERTa-v3-large models.

7 Discussion

We organize this discussion around the two research questions posed in Section 1.

RQ1: Does a fine-tuned math-specialized encoder match the performance of a larger general-purpose encoder on math-explanation items? Yes, if both are fine-tuned under an identical training harness. On the 18 items run under both encoders with identical splits, mean best QWK differed by only 0.002 (0.931 for MathBERT versus 0.933 for DeBERTa-v3-large), the models produced the best performance at similar rates (8 for MathBERT to 10 for DeBERTa-v3-large), and the models performed quite similarly across the other three acceptance criteria metrics (exact agreement, SMD, and scorepoint distribution). We therefore

suggest an approach that starts with fine-tuning the smaller, more efficient, MathBERT encoder that has been pre-trained specifically to interpret text that includes mathematical notation.

RQ2: Does escalating to a larger general-purpose encoder for items not meeting acceptance criteria improve performance? Yes. Of the six items for which a fine-tuned MathBERT encoder failed to meet acceptance criteria, a fine-tuned DeBERTa-v3-large model was able to flip the results to acceptable. This two stage modeling approach therefore worked for all but two of the items where human-to-human agreement was either abnormally high (0.949) or low (0.822), suggesting that the items may be outliers that are not well suited for automated scoring. Table 3 reports the exact agreement performance metrics for these six items.

8 Limitations

One limitation is the scope of this research. The results reported here comes from items modeled from a single administration of a single state-wide assessment. The results accordingly could lack external validity because item styles, rubric conventions, and rater training vary across programs.

Secondly, automated scores often influence high-stakes decisions, so agreement with human raters is necessary but not always sufficient. An engine must also score comparably well across student subgroups. Our analysis establishes aggregate human-like agreement but does not dis-aggregate by subgroup. That analysis should be conducted before any operational use and is an important item of remaining work.

9 Conclusion

Both a math-specialized encoder (MathBERT) and a larger general-purpose encoder (DeBERTa-v3-large) can be fine-tuned to human-like agreement on operational math-explanation items, and their respective performance is practically indistinguishable. MathBERT achieves the same performance at roughly 40% fewer parameters, allowing it to be fine-tuned more efficiently, and has been pre-trained specifically to interpret text that includes mathematical notations. We accordinglyi recommend it as the default pre-trained encoder to score math-explanation items, and to fine-tune DeBERTa-v3-large models only for items for which a fine-tuned MathBERT model is unable to meet operational deployment acceptance criteria.

References

Randy Elliot Bennett and Mo Zhang. 2011. Automated scoring of constructed-response literacy and mathematics items. Technical report, Educational Testing Service.

Jacob Cohen. 1968. Weighted kappa: Nominal scale agreement provision for scaled disagreement or partial credit. *Psychological Bulletin*, 70(4):213–220.

Scott Crossley, Perpetual Baffour, Tian Yu, Alex Franklin, Meg Benner, and Ulrich Boser. 2023. A large-scale corpus for assessing written argumentation: PERSUADE 2.0.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4171–4186.

Pengcheng He, Jianfeng Gao, and Weizhu Chen. 2023. DeBERTaV3: Improving DeBERTa using ELECTRA-style pre-training with gradient-disentangled embedding sharing. In *International Conference on Learning Representations*.

Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. DeBERTa: Decoding-enhanced BERT with disentangled attention. In *International Conference on Learning Representations*.

Scott Hellman, Alejandro Andrade, and Kyle Habermehl. 2023. Scalable and explainable automated scoring for open-ended constructed response math word problems. In *Proceedings of the 18th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2023)*, pages 137–147. Association for Computational Linguistics.

Ilya Loshchilov and Frank Hutter. 2019. Decoupled weight decay regularization. In *International Conference on Learning Representations*.

Wesley Morris, Scott Crossley, Langdon Holmes, and Amber Trumbore. 2025a. Mean squared error versus cross-entropy loss for ordinal automated scoring of constructed responses. *Journal of Educational Data Mining*.

William Morris, Langdon Holmes, Jinsook S. Choi, and Scott Crossley. 2025b. Automated scoring of constructed response items in math assessment using large language models. *International Journal of Artificial Intelligence in Education*, 35:559–586.

NAEP Research and Development Program. 2023. NAEP math automated scoring challenge results and preprint.

Pearson and ETS. 2015. Research results of parcc automated scoring proof of concept study. Technical report, Pearson and ETS. Accessed September 2026.

Jia Tracy Shen, Michiharu Yamashita, Ethan Prihar, Neil Heffernan, Xintao Wu, Ben Graff, and Dongwon Lee. 2021. MathBERT: A pre-trained language model for general NLP tasks in mathematics education. *arXiv preprint arXiv:2106.07340*.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30.

David M. Williamson, Xiaoming Xi, and F. Jay Breyer. 2012. A framework for evaluation and use of automated scoring. *Educational Measurement: Issues and Practice*, 31(1):2–13.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, et al. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45. Association for Computational Linguistics.

Duanli Yan, André A. Rupp, and Peter W. Foltz. 2020. *Handbook of Automated Scoring: Theory into Practice*. Chapman and Hall/CRC.

How Much Does Hyperparameter Tuning Actually Help? An Efficiency Survey for Fine-Tuning Transformers to Score Mathematics-Explanation Items

Gregory M. Jacobs **Ahmed H. Bediwy** **Martha Bellows**
Pearson Education
{greg.jacobs, ahmed.bediwy, martha.bellows}@pearson.com

Abstract

Fine-tuning pre-trained transformer models for constructed-response items often begins with a hyperparameter grid search using k-fold cross-validation. We study how much that search actually helps by fine-tuning two encoders—MathBERT, a smaller math-focused model, and DeBERTa-v3-large, a larger general-purpose model—to score 18 math-explanation items from a state-wide assessment. Crossing learning rate, weight decay, and label smoothing over two epochs and five folds (1,800 model-fold runs), we compare each model's tuning gain directly against fold-to-fold noise via a gain-to-noise ratio. Gains were small relative to fold noise for both encoders: MathBERT's ratio fell below one (0.90), and DeBERTa's nominally higher ratio (1.42) traced to a handful of divergent fits rather than an informative search landscape. Weight decay and label smoothing were effectively inert, leaving learning rate as the only hyperparameter worth checking—though even for learning rate the best configurations offered small performance gains over a reasonable default. We accordingly recommend a lean workflow that fixes the inert hyperparameters to sensible defaults, runs a narrow learning-rate search extended modestly upward, and increases the epoch budget with early stopping, substantially reducing training compute without sacrificing accuracy.

1 Introduction

Fine-tuning a pre-trained transformer model for scoring a constructed-response item conventionally begins with hyperparameter (HP) tuning. A common approach to HP is a grid search wrapped in k-fold cross-validation (Morris et al., 2025). The approach is so standard that it is rarely scrutinized. Most experiments simply enumerate a grid across a set of parameters, evaluate each configuration by an appropriate cross-validated metric, and keep the highest performing parameter values. For an operational program that includes dozens of mathematics-explanation items at multiple grades for every administration cycle, this search can be quite resource intensive because the number of training runs scales at a rate of items $\times$ configurations $\times$ folds. For example, even a modest grid of a dozen configurations over five folds already encompasses sixty fine-tuning runs *per item* before a single production model is trained.

Yet the assumption that exhaustive tuning is *necessary* is seldom scrutinized. One of the main challenges of this process is that each cross-validated QWK estimate is calculated on finite validation folds and therefore carries its own sampling noise. Refit the same configuration on different folds and the estimate often moves. When the apparent benefit of searching (the gap between the best configuration and a reasonable default) is smaller than this run-to-run variability, the search is effectively selecting noise. It will report a "winner," but that winner is not reliably reproducible and uncovers no dependable operational advantage. Quantifying this tradeoff could be significant when designing a fine-tuning pipeline that is both accurate and efficient at scale for a state-wide assessment program.

Conversely, if a particular model hyperparameter value genuinely needs tuning, a program must know which one so that it does not implement an overly lean pipeline that can result in degraded models. The goal therefore is to measure, per model and per hyperparameter, *when* tuning helps and *when* it does not, and to turn that into a concrete recipe.

This paper surveys how effective HP tuning actually is for automated scoring of a set of math-explanation items, using two encoders: MathBERT, a smaller, math-focused pretrained model (Shen et al., 2021), and DeBERTa-v3-large, a larger general-purpose model (He et al., 2021, 2023). Through this experimentation, we seek to answer the following research questions: (RQ1) Which hy-

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 200–206
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

perparameters, if any, improve performance when fine-tuning a smaller, math-focused encoder or a larger, general-purpose encoder to score math-explanation items, and (RQ2) what approach to hyperparameter tuning, if any, should be implemented when designing a training pipeline to fine-tune encoders for performing operational scoring of math-explanation items at scale?

2 Related Work

Hyperparameter optimization. The theory of hyperparameter optimization is well-developed. Bergstra and Bengio (2012) showed that random search is often as effective as grid search and far cheaper because only a few hyperparameters typically matter. Our study is complementary and specific to automated scoring. Rather than proposing a better search strategy, we ask whether searching is worthwhile at all once the resulting gains are compared against the sampling noise of the very metric used to select configurations.

Stability of fine-tuning. Parallel literature documents that fine-tuning large pretrained encoders can be unstable, with substantial variance across random seeds, data orders, and learning rates (Dodge et al., 2020; Mosbach et al., 2021). This instability is precisely the fold noise our diagnostics seek to quantify. If run-to-run variability is large relative to between-configuration differences, a grid search will ineffectively chase variance. We connect this general finding to the operational scoring setting and to a specific observed failure—the occasional divergence of the larger encoder at higher learning rates.

Cross-validation variance. Selecting the configuration with the highest cross-validated score amounts to maximizing over noisy performance estimates, making the selected configuration's apparent advantage susceptible to selection-induced optimism (Cawley and Talbot, 2010; Varma and Simon, 2006). The problem becomes more pronounced when the model-selection criterion has substantial variance, since model selection can itself overfit the cross-validation criterion (Cawley and Talbot, 2010). Moreover, k-fold cross-validation has nontrivial variance arising from the dependence among folds, and this uncertainty is not generally estimable without bias (Bengio and Grandvalet, 2004). Our gain-to-noise ratio provides a direct diagnostic of whether the observed tuning gain is large relative to this estimation noise. When the tuning gain is at or below the fold-noise floor, the "winner" may, to a first approximation, be the configuration that benefited most from validation noise rather than the configuration that generalizes best. Thus, the methodological question shifts from "which configuration wins?" to "is the observed win distinguishable from cross-validation noise?".

Automated scoring and evaluation. Automated scoring is evaluated against human-to-human reliability using ordinal agreement statistics (Cohen, 1968; Pearson and ETS, 2015; Williamson et al., 2012; Yan et al., 2020), and the 2023 NAEP Mathematics Automated Scoring Challenge showed that fine-tuned transformers can reach human-like agreement on such items (Crossley et al., 2023; NAEP Research and Development Program, 2023). Those systems emphasized peak accuracy through elaborate augmentation and ensembling. We instead study the *efficiency* of the tuning stage for operational deployment, treating compute as a constraint.

The structure of fine-tuning models for operational scoring. The resources required to fine-tune automated scoring models scale with the number of training runs, which a hyperparameter grid search multiplies by configurations and folds. Because these runs can reoccur every administration and across many items, even a constant-factor reduction compounds. Prior work on scoring accuracy sometimes fails to account for this, but we instead place it on the same compute-versus-accuracy footing as the rest of the pipeline.

3 Corpus

3.1 Data

Every statistic in this paper is computed from a single administration of 18 math-explanation items from one statewide summative assessment. The items span from grades 3 through 8 and apply rubrics with scorepoint ranges of either four (0-3) or five (0-4) integer values.

3.2 Human scoring

Responses were scored under an operational double-read procedure whereby 10% of responses received a second score to measure human-to-human reliability. For both training and evaluation, we take the target label to be `best_human`: the program's final reported score, determined by a

predefined hierarchy in which an adjudication or resolution score takes precedence over the first and second reads. Human–human reliability is taken from the first and second reads on the subset of responses that received a second human score.

3.3 Response pre-processing

Operational responses arrive in a semi-structured delivery format rather than as clean text, and they carry markup that a language model should not see verbatim. Before fine-tuning, we apply a deterministic cleaning pipeline: we unwrap the JSON delivery format, decode HTML entities and strip presentational tags, convert embedded MathML into a plain-text rendering so that symbols, fractions, and operators survive as readable tokens, and normalize Unicode. The pipeline is identical across all items and both encoders, so no difference between models can be attributed to pre-processing.

4 Hyperparameter Optimization Methodology

4.1 Broad Hyperparameter Search Grid.

The exploratory survey varied learning rate, weight decay, and label-smoothing α (Table 1). MathBERT crossed two values of each hyperparameter (8 configurations); DeBERTa-v3-large used three rates and two values for the other hyperparameters (12). Each configuration received a two-epoch budget and was evaluated by five-fold stratified cross-validation, yielding 1,800 model-fold runs (720 for MathBERT and 1,080 for DeBERTa). Because we are modeling math explanation items, the hyperparameters chosen to survey, and the breadth of their values included in the grid search, were based on the experimentation reported out from the 2023 NAEP Mathematics Automated Scoring Challenge (Crossley et al., 2023; NAEP Research and Development Program, 2023). Table 1 outlines the grid search values used for each encoder.

For each hyperparameter configuration, an item's 5,000 response training set was randomly divided into five 1,000 response folds for cross-validation. For each configuration/fold combination, a model is trained on four folds and validated on the held-out fold, using QWK as the validation metric. We then calculated the mean QWK value across the five folds. Figure 1 visually displays this five-fold cross-validation process.

Hyperparameter	Values searched
Learning rate	1e-5, 2e-5 (MathBERT); 1e-5, 1.5e-5, 2e-5 (DeBERTa)
Weight decay	0.01, 0.05 (MathBERT); 0.1, 0.3 (DeBERTa)
Smoothing α	0.5, 0.7
Batch size	4 (grad. accumulation)
Epochs	2
Scheduler	linear with warmup

Table 1: Hyperparameter grid search space. Grids yielded 8 (MathBERT) and 12 (DeBERTa-v3-large) configurations over 5 folds.

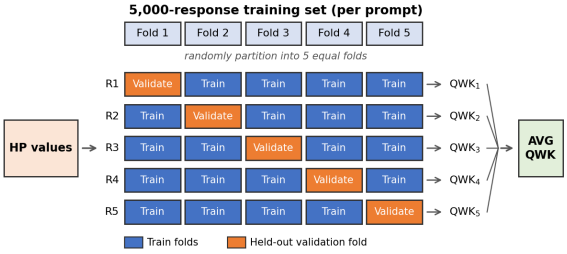


Figure 1: Five-fold cross-validated grid search. For a given hyperparameter configuration, the training set is partitioned into five folds; the model is trained on four folds and validated on the held-out fold, and the five QWK scores are averaged. This is repeated for every configuration in the grid.

4.2 Evaluation

For every prompt $\times$ model $\times$ configuration we compute the mean and standard deviation (SD) of validation QWK across the five folds (Figure 1). From these we derive four diagnostics.

Tuning gain the best-configuration mean QWK minus the average QWK across all configurations. This is what a search buys relative to picking a configuration at random from the grid.

Tuning range the best minus the worst configuration mean QWK. This is the most a search could possibly buy, i.e. the spread the grid exposes.

Fold noise the SD of the best configuration's QWK across its five folds. This is ordinary run-to-run variability at fixed hyperparameters.

Gain-to-noise ratio tuning gain divided by fold noise. A ratio below 1.0 means the benefit of searching is smaller than ordinary run-to-run variation and is therefore not reliably reproducible.

Formally, configurations by $c \in \{1, \ldots, C\}$ and folds by $j \in \{1, \ldots, 5\}$, and let q_{cj} be the validation QWK of configuration c on fold j. Write $\bar{q}_c = \frac{1}{5}\sum_j q_{cj}$ for a configuration's mean and $c^\star = \arg\max_c \bar{q}_c$ for the selected configuration. Then, for a prompt,

$$\text{gain} = \bar{q}_{c^\star} - \frac{1}{C}\sum_c \bar{q}_c, \quad (1)$$

$$\text{range} = \max_c \bar{q}_c - \min_c \bar{q}_c, \quad (2)$$

$$\text{noise} = \left(\frac{1}{4}\sum_j (q_{c^\star j} - \bar{q}_{c^\star})^2\right)^{1/2}, \quad (3)$$

$$\text{gain-to-noise} = \text{gain}/\text{noise}. \quad (4)$$

We report these per prompt and average across prompts within a model.

For each individual hyperparameter we additionally tabulate (i) how often each grid value "wins" (is chosen in the best configuration), (ii) its marginal mean QWK averaged over the other knobs, and (iii) whether the winning value sits at a grid edge. A winning value at a grid edge suggests that the grid might benefit from being *extended* rather than that tuning has converged. Finally, we examine the number of training epochs by aggregating each model's validation-QWK-versus-epoch curve.

Gain versus range. It is worth distinguishing the tuning *gain* from the tuning *range*. The range (best minus worst) can look large only because the grid includes a few genuinely bad configurations. But the relevant quantity is the gain over a *sensible* default, not the worst, which is what the gain diagnostic measures. A wide range with a small gain is the signature of a grid whose spread is skewed by avoidable bad settings rather than by meaningfully different good ones.

5 Results

Table 2 reports the tuning effectiveness by model. Across models the tuning gains were small and mostly of the same magnitude as fold noise. Table 3 reports the percentage each hyperparameter value resulted in the best performing configuration for each model. These results are analyzed for each model in the following two subsections.

5.1 MathBERT is robust to its hyperparameters

MathBERT was strikingly insensitive to the grid search parameters. The mean best QWK was 0.928, but the tuning gain was only 0.0053 against fold noise of 0.0058. With a resulting gain-to-noise ratio of only 0.90, the average gain of searching the grid was *smaller* than the variation from rerunning the best configuration on different folds.

Model	Mean best QWK	Tuning gain	Tuning range	Fold noise	Gain/ noise
MathBERT	0.928	0.0053	0.0107	0.0058	0.90
DeBERTa-v3-large	0.931	0.0106	0.0629	0.0075	1.42

Table 2: Tuning effectiveness by model (grid search, validation QWK). A gain-to-noise ratio below 1.0 means the search bought less than ordinary run-to-run variation.

Model	Hyperparameter	Winning value	Win rate
MathBERT	Learning rate	2e-5 (grid max)	100%
	Weight decay	0.01/0.05	0.61/0.39
	Smoothing α	0.5	67%
DeBERTa	Learning rate	2e-5 (grid max)	78%
	Weight decay	0.1/0.3	no pref.
	Smoothing α	0.5	56%

Table 3: Per-hyperparameter win rates. For DeBERTa-v3-large the winning learning rate sits at the top of the searched range, but the marginal differences across values are within fold noise.

Learning rate 2e-5 won on all 18 prompts, but with an edge of only about 0.007 QWK over the lower rate. So, even the "wrong" learning rate cost almost nothing. Weight decay was effectively inert: the two grid values split wins $0.61/0.39$ with a negligible marginal difference, indicating no reliable preference. Label-smoothing $\alpha = 0.5$ modestly outperformed $\alpha = 0.7$ (winning 67% of prompts). Taken together with the sub-1.0 gain-to-noise ratio, the per-parameter picture is that no MathBERT hyperparameter reliably matters within the searched ranges. A single sensible default performs as well as the grid winner once fold noise is accounted for.

5.2 DeBERTa-v3-large is subject to occasional divergence

For DeBERTa-v3-large, the gain-to-noise ratio (1.42) in the aggregate indicates significance, but the per-parameter details show that this might not be the case. The marginal mean QWK of the three learning rates differs by less than 0.004—smaller than the fold noise—so no learning-rate value is reliably better on average. What inflates the tuning *range* (and consequently the gain over the cross-configuration average) is a small number of unstable fits: 6 of the 216 DeBERTa configurations collapsed to a QWK between 0.67 and 0.79, against roughly 0.92 for the rest. These collapses are not

tied to a single learning rate hyperparameter value. They occur at 1.5e-5 and 2e-5 but never at the lowest rate (1e-5). This is consistent with the known instability of fine-tuning large encoders (Mosbach et al., 2021; Dodge et al., 2020). DeBERTa-v3-large's favorable ratio is therefore evidence that one must always be cognizant of the occasional divergent fit when fine-tuning a larger encoder model.

Interestingly, the winning learning rate sat at the top of the searched range (2e-5) for 78% of prompts, yet the divergences also occurred at the two higher rates. The best runs accordingly occurred at the slightly higher rates, but that same region carries the divergence risk. The sensible response therefore is to keep the learning rate in the conservative 1–2e-5 band and retain a one- or two-point check to catch the rare collapse rather than to refine an interior grid.

6 Discussion

Our first research question asks which hyperparameters, if any, improve performance when fine-tuning a smaller, math-focused encoder or a larger, general-purpose encoder to score math-explanation items. We address this question, organized by hyperparameter, below.

6.1 Learning Rate

As noted in the previous section, tuning the learning rate hyperparameter on the grid values used in the experiment did not meaningfully impact model performance in the majority of circumstances for either encoder (with the exception of resulting in a rare divergence for DeBERTa models). It is noteworthy, however, that the best performing MathBERT models used the grid's maximum learning rate value (2e-5) for every model. The best performing DeBERTa models similarly used the grid's maximum learning rate value (2e-5) a majority of the time (78%). This suggests that it may be prudent to extend the grid search values upward for any future hyperparameter tuning of these models for math explanation items, even if it is unlikely to significantly improve performance over the default values.

6.2 Weight decay and label smoothing

Neither the weight decay or label smoothing behaved like a lever worth tuning. Weight decay was effectively inert for both encoders, splitting wins near evenly with a marginal effect negligible relative to fold noise. Weight decay resulted in a slight preference of 0.01 (61%) for MathBERT models and no preference between 0.1 and 0.3 for DeBERTa models. Label-smoothing α showed a mild, consistent preference for 0.5 over 0.7 for both models. These results indicate that both hyperparameter values can be *fixed* rather than searched in future hyperparameter tuning of encoder models for scoring math explanation items.

6.3 Number of training epochs

Both models show significant improvement between the first and second training epochs. MathBERT mean validation climbed from 0.889 at one epoch to 0.928 at two epochs. DeBERTa-v3-large rises from 0.878 at one epoch to 0.931 at two epochs. This increase between one and two epochs suggests that increasing the number of epochs beyond 2 could bear additional gains, though there does not appear to be much more to gain at mean QWK values exceeding 0.9 for both models.

7 Recommendations

Our second research question asks what approach to hyperparameter tuning, if any, should be implemented when designing a training pipeline to fine-tune encoders for performing operational scoring of math-explanation items at scale. The experimentation results support a workflow that includes a targeted learning rate grid search as the sole hyperparameter tuned for both smaller math-specific and larger general purpose encoders. We recommend the following model training approach when fine-tuning theses models for scoring math-explanation items:

1. **Fix all non-learning rate hyperparameter values.** Set reasonable default values for weight-decay, smoothing, scheduler, and batch settings rather than including them in any hyperparameter grid search. Based on the winning values observed in this study, we suggest the following defaults:

    - **Weight decay:** 0.01 for MathBERT (won 61% of prompts); 0.1 for DeBERTa-v3-large (the two grid values were effectively tied, so the more conservative value is retained).

    - **Label smoothing α:** 0.5 for both encoders (consistently preferred over 0.7).

- **Scheduler:** linear schedule with warmup for both encoders.

- **Batch size:** 4 with gradient accumulation of 4 for both encoders, resulting in an effective batch size of 16.

2. **Implement informed learning rate grid search.** Because the best configurations for both encoders overwhelmingly favored the top of the range of grid search values for learning rate (2e-5), we suggest the following targeted learning-rate grids for future trainings:

   - **MathBERT:** $\{2, 3, 4, 5\}$e-5, extending the grid search to capture any marginal gains that may exist at the higer rates, and

   - **DeBERTa-v3-large:** $\{1.5, 2, 3, 4\}$e-5, kept in the conservative band to capture the best runs while retaining a low-rate anchor that guards against the occasional divergent fit at higher rates.

3. **Increase number of epochs with early stopping.** Because validation QWK increased between one and two epoch trainings for both encoders, we recommend increasing the number of training epochs, particularly for the smaller MathBERT encoder, while using a more conservative increase for the larger DeBERTa-v3-large encoder. We also recommend implementing early stopping with the expectation that results could plateaus within a few additional epochs given the large validation QWK values produced after two training epochs.

Because search cost scales as items $\times$ configurations $\times$ folds $\times$ training epochs, streamlining the hyperparameter tuning approach as recommended above will result in a significant savings in training time and resources.

8 Limitations

This experiment includes a group of 18 constructed response math items from one administration, so the results are generalizable within a caveat of a fairly limited sample size and subject matter. Formal equivalence tests and cross-administration replication remain for future investigative work on this research.

9 Conclusion

This research surveyed how much hyperparameter tuning actually helps when fine-tuning MathBERT and DeBERTa-v3-large encoders to score math-explanation items. Across the 18 items, tuning gains were small relative to fold noise for both encoders: MathBERT's gain-to-noise ratio fell below one (0.90), and DeBERTa's nominally higher ratio (1.42) traced to mainly to a handful of divergent fits rather than an informative search landscape. Weight decay and label smoothing were effectively inert and can be fixed to sensible defaults, leaving learning rate as the only hyperparameter worth tuning using an informed and targeted grid search, though even for learning rate there the best configurations offered small performance gains over a reasonable default.

These findings support a lean, compute-efficient hyperparameter tuning approach: fix the non-learning-rate hyperparameters to defaults, run a narrow learning-rate grid search extended modestly upward, and increase the epoch budget with early stopping to capture any additional gains that may be attainable by training beyond two epochs. Because tuning cost scales with items, configurations, folds, and epochs, this streamlined approach substantially reduces training time and compute while preserving the small, prompt-specific benefits that tuning could provide. Formal equivalence testing and cross-administration replication remain promising directions for confirming these conclusions.

References

Yoshua Bengio and Yves Grandvalet. 2004. No unbiased estimator of the variance of k-fold cross-validation. *Journal of Machine Learning Research*, 5:1089–1105.

James Bergstra and Yoshua Bengio. 2012. Random search for hyper-parameter optimization. *Journal of Machine Learning Research*, 13:281–305.

Gavin C. Cawley and Nicola L. C. Talbot. 2010. On over-fitting in model selection and subsequent selection bias in performance evaluation. *Journal of Machine Learning Research*, 11(70):2079–2107.

Jacob Cohen. 1968. Weighted kappa: Nominal scale agreement provision for scaled disagreement or partial credit. *Psychological Bulletin*, 70(4):213–220.

Scott Crossley, Perpetual Baffour, Tian Yu, Alex Franklin, Meg Benner, and Ulrich Boser. 2023. A large-scale corpus for assessing written argumentation: PERSUADE 2.0.

Jesse Dodge, Gabriel Ilharco, Roy Schwartz, Ali Farhadi, Hannaneh Hajishirzi, and Noah A. Smith.

2020. Fine-tuning pretrained language models: Weight initializations, data orders, and early stopping.

Pengcheng He, Jianfeng Gao, and Weizhu Chen. 2023. DeBERTaV3: Improving DeBERTa using ELECTRA-style pre-training with gradient-disentangled embedding sharing. In *International Conference on Learning Representations (ICLR)*.

Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. DeBERTa: Decoding-enhanced BERT with disentangled attention. In *International Conference on Learning Representations (ICLR)*.

Wesley Morris, Langdon Holmes, Joon Suh Choi, and Scott Crossley. 2025. Automated scoring of constructed response items in math assessment using large language models. *International Journal of Artificial Intelligence in Education*, 35(2):559–586.

Marius Mosbach, Maksym Andriushchenko, and Dietrich Klakow. 2021. On the stability of fine-tuning BERT: Misconceptions, explanations, and strong baselines. In *International Conference on Learning Representations (ICLR)*.

NAEP Research and Development Program. 2023. NAEP math automated scoring challenge results and preprint.

Pearson and ETS. 2015. Research results of parcc automated scoring proof of concept study. Technical report, Pearson and ETS. Accessed September 2026.

Jia Tracy Shen, Michiharu Yamashita, Ethan Prihar, Neil Heffernan, Xintao Wu, Ben Graff, and Dongwon Lee. 2021. MathBERT: A pre-trained language model for general NLP tasks in mathematics education. *arXiv preprint arXiv:2106.07340*.

Sudhir Varma and Richard Simon. 2006. Bias in error estimation when using cross-validation for model selection. *BMC Bioinformatics*, 7(1):91.

David M. Williamson, Xiaoming Xi, and F. Jay Breyer. 2012. A framework for evaluation and use of automated scoring. *Educational Measurement: Issues and Practice*, 31(1):2–13.

Duanli Yan, André A. Rupp, and Peter W. Foltz. 2020. Handbook of automated scoring: Theory into practice. *Chapman and Hall/CRC*.

Responsible Use of Generative AI When Creating Reading Comprehension Questions: Inference Matters

Zuowei Wang, Michael Flor, Jiyun Zu, Tenaha O'Reilly, Wanjing Anya Ma
Educational Testing Service
zwang@ets.org, mflor@ets.org, jzu@ets.org, toreilly@ets.org, anyawma08@gmail.com

Abstract

This work-in-progress paper shows how traditional item-quality indicators can miss construct-relevant differences between AI-generated and expert-created multiple-choice (MC) reading comprehension (RC) items. Items from both sources passed quality review by independent expert reviewers, and responses from 148,825 test takers showed similar item difficulty, item slope, and item fit. However, further analysis revealed differences in the distribution of required inference types. These discrepancies may escape traditional quality review but can be detected by examining whether the AI model faithfully followed prompts during intermediate item-generation steps. Evaluations of AI-generated items should therefore document prompts and generation steps to help identify and mitigate construct-relevant differences.

1 Introduction

Multiple-choice (MC) reading comprehension (RC) questions have been a staple measure of students' academic achievement for educators and researchers. The creation of high-quality MC RC questions requires expertise and effort. Recently, large language models (LLMs) have been used to generate RC items (Flor, 2025). Empirical studies have demonstrated that MC RC questions created by LLMs can have similar properties to those by human experts based on traditional indicators (Shin et al., 2025; Tomikawa & Uto, 2024). As more assessment providers use LLMs to create MC RC questions, a critical question is whether traditional quality indicators are sufficient to ensure the quality of these items.

This work-in-progress paper compares AI-generated and expert-created multiple-choice reading comprehension (MC RC) items that had all passed the same standardized quality review process. We first examined whether the two item sources differed on traditional quality indicators, including item slope, item difficulty, item fit, and the relation between passage text complexity (Sheehan et al., 2014) and item difficulty. We then compared them on an additional dimension central to the construct of reading comprehension: the types of inferences required by the items (Cain & Oakhill, 1999; Graesser et al., 1994; Singer et al., 1992). The goal was to determine whether AI-generated and expert-created items that appeared comparable by traditional quality indicators could still differ in the inference processes they required. Because such differences may not be detected by traditional quality indicators, we argue that evaluations of AI-generated items should document the prompts and generation steps used to create them. In the session, we present a proposed approach and invite feedback on its potential to help ensure that AI-generated items function as intended.

2 Method

2.1 Sample

The study sample included 148,825 students in grades 3–12 from the United States who completed a standardized RC test in 2025. Schools used this

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 207–210
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

test to identify struggling readers and inform instruction. Periodically, the test includes a trial passage and associated items so that item parameters can be estimated before the passages become operational. Student responses to trial-passage items do not contribute to final scores, and students are not told which passage is being trialed.

2.2 Item creation

We edited 20 public-domain passages to be used as trial passages, so they had an average text complexity of grade 7 (SD = 2 grade levels; Sheehan et al., 2014). We randomly divided the passages into two groups of 10. For the first group, professional assessment developers wrote MC RC questions for each passage. For the second group, we used GPT-4o to generate MC RC questions using methods reported in Ma et al. (2025). Specifically, we used few-shot examples to create four types of RC MC items: gap-filling inference, text-connecting inference, pronominal bridging inference, and literal comprehension, which required little inference making.

2.3 Item quality review

Questions from both groups underwent a standardized quality review process used for high-stakes tests. The review accepted 153 MC questions. These questions were dispersed in a standardized RC test to appropriate grade levels as trial items.

2.4 Item calibration and analysis of item parameters

In this study, each trial item received responses from 3,525 to 14,740 students in the sample. Because different trial passages were used in different grade levels, and because the trial passages were administered periodically, students only responded to a subset of the items. We calibrated the items using item response theory (IRT) with a 2PL model (Embretson & Reise, 2013). Two experts-created items and one AI-generated item had very high difficulty (i.e., b>3) or unusually large item slope (i.e., a>10) and were removed from subsequent analysis, since they would not be included in the operational item pool after the standard item parameter review protocol of the test. This resulted in 71 AI-generated and 79 expert-created MC questions.

We compared the item parameters and item fit statistics of the two sets of remaining questions. We also examined whether text complexity was similarly related to the difficulty of the items in both question sets.

2.5 Manual coding of items' inference types

Two human coders classified each question based on the type of inference it required (i.e., gap-filling, pronominal bridging, text-connecting, and literal). To develop a shared understanding of the coding scheme, the coders first jointly coded and discussed 18 expert-created questions from two passages. They then independently coded the remaining 142 questions from 18 passages. These independent ratings were used to calculate inter-rater agreement.

Nine expert-created MC questions did not fit the four categories because they asked students to determine the meaning of specific words in context. Both coders labeled these items as vocabulary-in-context questions. Because AI was not used to generate this item type in the present study, these questions were excluded only in the inter-rater agreement analysis and from the comparison of inference processes across the two question sets.

We examined inter-rater agreement and then compared the distribution of inference-question types across the two item sources. Further, we examined the few-shot prompts used to generate the questions and analyzed to what extent the generated questions conformed with the corresponding prompt.

3 Results

3.1 AI-generated and expert-created items show similar item parameters.

The expert-created and AI-generated MC questions had similar item parameters and item fit statistics, including slope, difficulty, and RMSEA – we found no significant difference between expert-created and AI-generated items on slope, difficulty, or RMSEA. In both sets of items, passage text complexity (Sheehan et al., 2014) predicted item difficulty in a similar way; there was no interaction between item source and text complexity. Table 1 shows the detailed results.

	Expert	GPT-4o
N (items)	79	71
N (passages)	10	10
Mean (Text Complexity)	707	706
SD (Text Complexity)	205	172
Range (Slope)	[.36, 2.59]	[.36, 2.10]
Mean (Slope)	1.31	1.26
SD (Slope)	0.44	0.44
Range (Difficulty)	[-.82, 2.33]	[-.86, 1.88]
Mean (Difficulty)	0.01	0.03
SD (Difficulty)	0.59	0.59
Range (RMSEA)	[0, 0.01]	[0, 0.02]
Mean (RMSEA)	0.006	0.006
SD (RMSEA)	0.003	0.003

Table 1: Distribution of Text Complexity and Item Parameters by Source

3.2 Inter-rater reliability when coding inference type of items

When coding the type of inference of items, the two human coders' inter-rater reliability as measured by Cohen's Kappa is .76 for the 52 independently coded expert-created items and .76 for the 71 independently coded AI-generated items, suggesting substantial agreement.

	AI-Generated	Expert-Created
Gap-Filling	9 (18%)	20 (42%)
Pronominal Bridging	30 (59%)	19 (40%)
Text-Connecting	12 (24%)	9 (19%)
Total	51 (100%)	48 (100%)

Table 3: Number of questions by inference type and source

3.3 AI-generated and expert-created items show different inference type distributions.

Forty-eight (61%) of the 79 expert-created items and 51 (72%) of the 71 AI-generated items required inference making. The manual coding revealed differences in the distribution of inference-question types across item sources. Specifically, AI tended to create more pronominal bridging and text-connecting questions but fewer gap-filling questions, $\chi^2(2) = 6.99$, $p < .05$. Table 2 shows the different distributions of inference types across item sources.

	Gap-Filling	Pronominal Bridging	Text-Connecting
Gap-filling	9	0	0
Pronominal Bridging	7	15	8
Text-Connecting	3	1	8
Literal	8	5	7

Table 2: Inference questions: target (columns) vs. actual (rows)

3.4 The AI model often failed to generate the target inference question specified by the prompt.

Further analysis of LLM's item generation process revealed that it often failed to produce specific types of questions as instructed in the prompt. For example, when the prompt asked the LLM to create gap-filling questions, only 9 of the 27 questions created were actually gap-filling questions, with 7 coded as pronominal bridging, 3 coded as text-connecting, and 8 as literal questions that did not require inference making (Table 3). Overall, only 45% of the questions generated by the LLM were consistent with the target inference type specified by the prompt. This failure to create specific types of inference questions appears to explain the different composition of item types in the expert-created vs. AI-generated questions.

4 Discussion

The standard item review process adopted by the RC test did not examine whether questions required the intended types of inference. If we had relied only on the usual quality review process, the difference in the distribution of inference types between expert-created and AI-generated MC RC questions would have remained hidden. Because inference making is central to reading comprehension, even subtle shifts in the kinds of inferences elicited by AI-generated items could affect what the test measures. The present findings

therefore suggest that AI-generated RC items may contain important construct-relevant differences that are not captured by traditional quality indicators alone.

This study shows that evaluating AI-generated RC questions requires looking beyond final item statistics and reviewer judgments. Here, the AI-generated items met predetermined standards and showed comparable psychometric properties, yet they did not consistently reflect the inference-question types specified in the prompt. This mismatch became visible only when we examined the intermediate generation process. We therefore argue that responsible evaluation of AI-generated RC questions should include transparent documentation of prompts, generation steps, and verifications of whether generated items conform to the intended design.

This kind of process-level evaluation also creates opportunities for mitigation. Once prompt-following failures are identified, developers can adjust prompts, add targeted review criteria, or intentionally generate additional underrepresented item types. For example, if LLMs underproduce valid gap-filling inference questions, developers could generate a larger pool of such questions and review them specifically for inference quality. In this way, transparency about the generation process is not only a safeguard against hidden weaknesses but also a practical path toward improving AI-facilitated item development.

References

Cain, K., & Oakhill, J. V. (1999). Inference making ability and its relation to comprehension failure in young children. *Reading and Writing, 11*(5-6), 489-503.

Embretson, S. E., & Reise, S. P. (2013). *Item response theory*. Psychology Press.

Flor, M. (2025). *Automatic question generation*. Springer.

Graesser, A. C., Singer, M., & Trabasso, T. (1994). Constructing inferences during narrative text comprehension. *Psychological review, 101*(3), 371. https://doi.org/10.1037/0033-295X.101.3.371

Ma, W. A., Flor, M., & Wang, Z. (2025). Automatic generation of inference making questions for reading comprehension assessments. *arXiv preprint arXiv:2506.08260*.

Sheehan, K. M., Kostin, I., Napolitano, D., & Flor, M. (2014). The textevaluator tool: Helping teachers and test developers select texts for use in instruction and assessment. *The Elementary School Journal, 115*(2), 184-209. https://doi.org/10.1086/678294

Shin, D., Kwon, S. K., & Lee, Y. (2025). Examining the efficacy of generative artificial intelligence in item generation: Comparative analysis of human-developed and ai-generated reading tests. *Education and Information Technologies, 30*(16), 23981-24007. 10.1007/s10639-025-13683-6

Singer, M., Andruslak, P., Reisdorf, P., & Black, N. L. (1992). Individual differences in bridging inference processes. *Memory & Cognition, 20*(5), 539-548. 10.3758/BF03199586

Tomikawa, Y., & Uto, M. (2024). Difficulty-controllable multiple-choice question generation for reading comprehension using item response theory. In A. M. Olney, I.-A. Chounta, Z. Liu, O. C. Santos, & I. I. Bittencourt, Artificial Intelligence in *Education. Posters and Late Breaking Results, Workshops and Tutorials, Industry and Innovation Tracks, Practitioners, Doctoral Consortium and Blue Sky Cham*.

Fine-tuning DeBERTa-v3 to Automate Spatial Language Classification

Sam Agnoli
Northwestern University

Qingzhou Shi
Northwestern University

David Uttal
Northwestern University

Abstract

To automate spatial language classification, we fine-tuned DeBERTa-v3, a transformer-based language model, using a 33,284-word dataset following a 70:15:15 train-validation-test split. The model performed well in a held-out test comparing performance to human coders (Cohen's kappa = .88).

Keywords: Language classification; spatial language; fine-tuning; educational measurement; large language models

Structured Summary

1 Background

Children's spatial skills, including the ability to mentally visualize and manipulate objects within their environment (Uttal et al., 2013), are predictive of success in STEM fields later in life (Lauer & Lourenco, 2016; Frick, 2019). One way to index the degree of spatial abilities in children is through the analysis of their spatial language in conversations (Fox et al., 2025; Pruden et al., 2011). Spatial language includes words for location and direction (e.g., in, up), dimension and size (e.g., small, tall, shallow), shape (e.g., circle, curved), and orientation or transformation (e.g., turn, flip). Determining whether language is spatial is difficult because many of these words are used spatially only in specific contexts. For example, to can express a spatial relation in "go to the top." but is non-spatial in "I want to help."; similarly, feet refers to a spatial measurement in "You are four feet tall." but to a body part in "My feet hurt." Because of the contextual complexity of language, researchers historically could not automate coding and instead relied on human coders, which remains an expensive and time-consuming bottleneck for research. Transcribing and coding open-ended language can take eight times as long as the recording itself (d'Apice et al., 2019; Margolin et al., 1998). Partly for this reason, spatial ability has more often been measured with non-language based spatial ability assessments, which are quick to administer and score despite being less ecologically valid (Harris, 2023).

Language models offer a scalable, cost-effective alternative to human coding. For example, recent work demonstrated that a multi-agent foundation model (i.e., GPT-4o with Administrator, Judger, and Critic agentic roles) performed similarly to human coders and identified novel spatial words (Shi & Uttal, 2025), which expanded the dictionary of possible spatial candidate words (Cannon et al., 2007). In the current study, we attempt to further improve the fidelity of language model coding by fine-tuning the DeBERTa-v3 base model on a large dataset of words previously coded by humans.

2 Aims

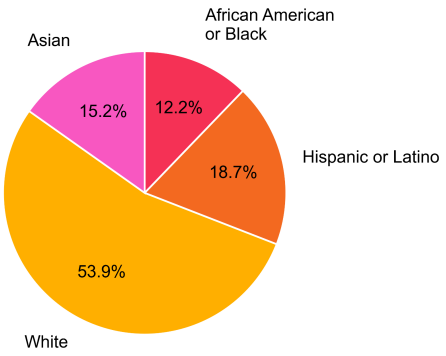
This study examined the validity of a fine-tuned DeBERTa-v3 model for spatial language coding (e.g., Yu et al., 2023). We compared a fine-tuned AI language model as a coder against traditional human coding, evaluated their performance, and explored implications for educational assessment.

3 Sample

Our sample consisted of recordings of 33,284 words across 239 parent-child conversations at a tinkering exhibit in the Chicago Children's Museum (see Polinsky et al., 2023). These counts reflect the final analytic sample after excluding eight sessions for which transcript and human-coding records could not be reliably matched due to inconsistencies between the transcript and coding files. Participating children ranged in age from 6 to 11 years, and families reflected the ethnic and socioeconomic diversity of Chicago (Figure 1). All families spoke fluent English.

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 211–215
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

Figure 1. *Participant Demographics*



The spatial dictionary (a predefined list of spatial candidate words that may have spatial meanings) developed by Cannon et al. (2007) was used to identify candidate words requiring contextual classification. The dictionary flagged 3,516 words as spatial candidates, representing 10.6% of all words and spanning 112 unique word types. Human coders subsequently judged whether each candidate was spatial in its utterance context. Approximately 41.4% of candidates were coded as spatial, corresponding to 4.4% of all words. The ten most frequent word types account for 61% of all spatial words.

4 Methods

Human coding followed the procedure detailed in Polinsky et al. (2023), in which a Bag-of-Words dictionary (Zhang et al., 2010) based on Cannon et al. (2007) flags candidate words that trained coders then judge in utterance context. This human-coded dataset serves as the criterion against which model performance is evaluated.

We used the open-source transformer-based language model DeBERTa-v3 base (He et al., 2021, 2023) for spatial language coding because of its efficiency in natural language understanding tasks. DeBERTa-v3 builds on DeBerta's disentangled attention, which separates a word's meaning from its position, making it well-suited for binary word-level classification (spatial = 1, non-spatial = 0) with the full utterance provided as context. Spatial words were identified using Polinsky et al. (2023).

The fine-tuned model was trained using a 70:15:15 train-validation-test split grouped by parent-child conversation to prevent data leakage (see Raschka, 2018). Given that spatial words are rare (4.4% of all words), we did not apply class weighting or resampling. Instead, the spatial dictionary restricted scoring to candidate words at inference, and checkpoint selection used validation minority-class F1 (computed over all words) rather than accuracy. We trained the model over two epochs to limit overfitting (see Devlin et al., 2019) and because that F1 did not improve on runs with three to five epochs. We then evaluated the held-out test set against human coding on both overall and candidate-only classification. After fine-tuning, we calibrated the model's decision confidence using temperature scaling (Guo et al., 2017), a standard method to reduce model overconfidence.

Performance of the fine-tuned model was evaluated using the following metrics: Cohen's kappa, accuracy, sensitivity (recall), specificity, positive predictive value (PPV or precision), negative predictive value (NPV), and F1 score. All metrics range from 0 to 1 except for Cohen's kappa, which ranges from -1 to 1. Higher values reflect stronger consistency with human coding. We also computed the error metrics expected calibration error (ECE; Guo et al., 2017) and Brier score (Brier, 1950) before and after temperature calibration. ECE summarizes the discrepancy between predicted confidence and observed accuracy across confidence bins, whereas the Brier score measures the squared difference between predicted probabilities and observed binary outcomes. For both ECE and Brier scores, lower values indicate better probabilistic predictions.

5 Results

We tested the fine-tuned DeBERTa-v3 model on a held-out dataset of 526 spatial-candidate words from a total of 4,470 words. The held-out dataset was created using a group-aware split by conversation, ensuring no conversation overlap between training and test data.

The performance metrics are presented in Table 1, with the positive class representing spatial words. The fine-tuned model achieved strong agreement with human coders across all words (i.e., overall evaluation set Cohen's kappa = .88) and within spatial candidate words (Cohen's kappa = .81). Because 95.3% of the test words are non-spatial, overall accuracy is inflated by class imbalance. The candidates-only evaluation is the more informative test, in which an accuracy of .91 compares against a .60 majority-class baseline. The fine-tuned model also produced well-calibrated probability estimates

Table 1. *Performance Metrics Across Fine-tuned DeBERTa-v3*

Accuracy		Sensitivity		Specificity		Precision (PPV)		NPV		F1	
FD	FD_s	FD	FD_s	FD	FD_s	FD	FD_s	FD	FD_s	FD	FD_s
0.99	0.91	0.92	0.92	0.99	0.90	0.84	0.86	1.00	0.95	0.88	0.89

Note. FD = Fine-tuned DeBERTa-v3; FD_s = Fine-tuned DeBERTa-v3 including spatial candidates only. The base rate of spatial words on the test set is 4.7% overall and 40.1% among spatial candidates. Accuracy should be interpreted against these prevalences.

Table 2. *Error Metrics of Fine-Tuned DeBERTa-v3*

Evaluation Set	Method	ECE Uniform	ECE Quantile	Brier	n (words)
Overall	Raw	0.008	0.004	0.010	4470
	Temp	0.005	0.003	0.009	
Spatial Candidates	Raw	0.059	0.044	0.074	526
	Temp	0.050	0.041	0.069	

Note. Raw = pre-calibration; Temp = post-hoc temperature calibration; ECE = expected calibration error (lower is better). ECE Uniform uses equal-width confidence bins; ECE Quantile uses equal-frequency bins. Brier = mean squared error of predicted probabilities (lower is better).

which further improved after post-hoc temperature calibration (see Table 2 and Figure 2). Taken together, these fine-tuning results are promising as a scalable alternative to human coding.

To provide calibrated confidence estimates for the model's decisions, we adjusted how closely the model's decision probabilities matched observed outcomes. Notably, calibration only changed what the model claimed to know, and left coding decisions untouched. The raw probabilities were overconfident within the spatial-candidate set (ECE = .059, Brier = .074). We fit a single temperature parameter on the candidate words of the validation split to address the model's overconfidence, yielding T = 1.46 (see Figure 2). Applying that temperature to the held-out test set reduced expected calibration error by roughly 15% (.059 to .050 under uniform binning; .044 to .041 under quantile binning) and reduced the Brier score from .074 to .069 (see Table 2). Metrics computed across all words were already near ceiling and changed little (ECE = .008 to .005), which is expected given that 95.3% of words are trivially classified as non-spatial. Calibration allows researchers to more readily accept high-confidence judgments and manually review cases where the model is less confident.

6 Conclusion

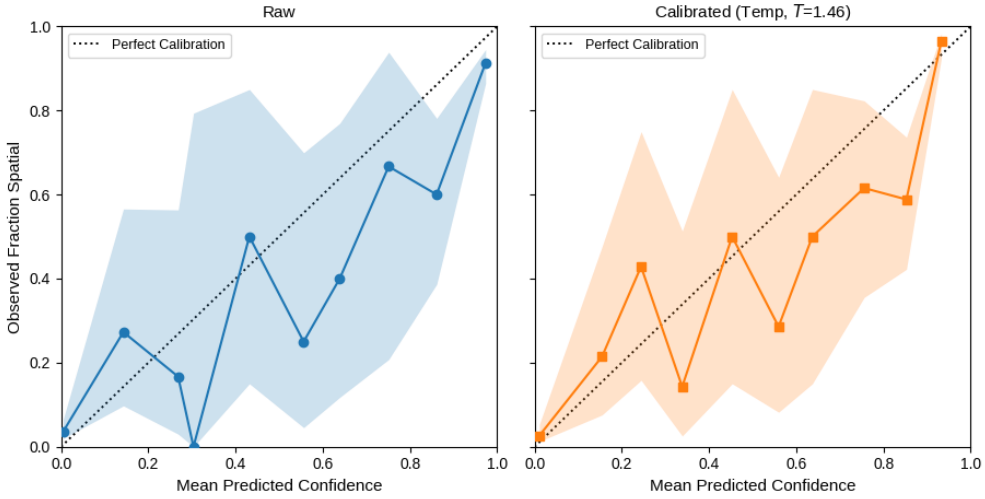
This study contributes to the field of AI in educational measurement by supplying an open-access (https://github.com/SamAgnoli/spatial-language-classifier/tree/main), scalable model for spatial coding. Our findings provide empirical evidence that a fine-tuned DeBERTa-v3 model can achieve strong agreement with human coders.

We suggest that researchers can deploy this model as a supplement to the costly process of human coding for spatial language. However, we caution against using this model as a replacement for human coding until the generalizability concerns outlined below are addressed. As an example, the current model may be sufficient for identifying relative differences or rank-ordering spatial ability within a homogenous sample. However, the model does not capture absolute values in spatial ability and should not be used as a diagnostic tool. We note that humans also make errors in spatial coding, and that the model may fall within an acceptable range of human variance for researchers' or educators' specific goals.

More generally, the value of this research may lie in what automated coding supports rather than the specific performance of a given model. Spatial ability has historically been measured with tractable, easy to score laboratory tasks (Harris, 2023). These tasks have been criticized for failing to capture real-world behaviors and abilities. In contrast, natural language offers a more ecologically valid and perhaps powerful tool to measure spatial ability. Whereas previously natural language was too cost- and time-intensive to analyze without ample resources, researchers can now record students in any situation of interest and be rewarded with a rich dataset. Automated coding may eventually overtake lab tasks in both availability and ecological validity, albeit with potential confounds (e.g., socioeconomic status, Noble et al., 2007).

There are several limitations to consider. We note that our data are limited (i.e., a few tasks in a museum tinkering exhibit) among English-speaking parent-child dyads. Also, the fine-tuned model was limited by its spatial dictionary. Future

Figure 2. *Calibrated Reliability Diagram for Spatial Candidate Subset*



Note. Figure 2 presents the reliability diagram for the spatial-candidate subset. Both line graphs track the diagonal (i.e., perfect calibration) at the extremes of the confidence range. The temperature calibrated line-graph is closer to the diagonal than the raw line-graph across most of that range, demonstrating better calibration. Shaded bands are 95% Wilson score confidence intervals for each bin's observed proportion, and are wide wherever a bin contains few words. Notably, 89.7% of candidate words fall in the lowest and highest confidence bins (n = 267 and n = 169, respectively, out of 526 candidates), so the volatility in the middle bins reflects greater noise due to the small sample.

research might develop a pipeline for expanding the dictionary for a given task following the multi-agentic approach outlined in Shi and Uttal (2025) and then fine-tuning DeBERTa-v3 (or a state-of-the-art model) provided researchers have sufficient data (i.e., spatial words within the utterance context). Additionally, human coding was based on the Cannon et al. (2007) and Polinsky et al. (2023) dictionary, whereas the model's candidate gate might be expanded following testing with the expanded dictionary from Shi and Uttal (2025). Re-coding the expanded candidates would allow the model's performance on novel spatial terms to be evaluated directly. Finally, to ensure generalizability, it is crucial to externally validate the model on independent datasets of varying demographics. Because the ten most frequent spatial words account for 61% of positive spatial language in the dataset and are domain-general (e.g., in, on, up, there), we suspect that the model's classification abilities would transfer to other (i.e., non-museum) English-speaking contexts where middle-childhood participants are tasked with describing spatial objects and relations. Regarding spatial classification in other languages with unique spatial representations (see Bowerman & Choi, 2001), researchers could apply our methodology to develop their own trained classifier, presuming they have sufficient human-coded spatial data.

Broadly, we underscore the importance of transparency in AI-related research as well as ensuring human alignment in AI-generated coding. This includes ensuring interpretable AI outputs, which we have attempted by using a spatial language dictionary and confidence calibration so that researchers can track which words or situations lead the model to be incorrect and/or overconfident. Such interpretability is essential for keeping humans aware and in control of model outputs. We also encourage other researchers to share their datasets either privately or in an open-source repository so that future researchers can use these datasets to train more advanced models as they become available.

Ultimately, this research advances the use of AI to support high-quality, context-aware educational measurement at scale.

7 Open Science and Code Availability

The fine-tuned model is publicly released under an MIT license at https://huggingface.co/SamAgnoli/deberta-v3-base-spatial-language-detection, and the full pipeline is available at https://github.com/SamAgnoli/spatial-language-classifier. The Hugging Face model card documents the input format, training hyperparameters, per-class evaluation metrics, and the temperature-scaling procedure, and the model can be loaded to classify

spatial language through a quick-start notebook accessible via the GitHub repository. In this repository, spatial words are identified using the dictionary developed by Polinsky et al. (2023), though we encourage future human coders to use the expanded dictionary by Shi and Uttal (2025). The repository also contains the annotated notebook covering dictionary construction and candidate matching, the conversation-grouped 70:15:15 split, fine-tuning, held-out evaluation on both the overall and candidates-only views, and post-hoc calibration.

Researchers applying the classifier should note that it is a sentence-pair classifier: it judges a target word together with its utterance. Performance outside English-language, parent-child tinkering speech has not been established, and we recommend the candidates-only metrics as the appropriate basis for judging model quality. Consistent with our recommendation that researchers share data where possible, we have provided a deidentified variant of our human-coded dataset within the repository under the data/example subfolder.

References

Bowerman, M., & Choi, S. (2001). Shaping meanings for language: Universal and language-specific in the acquisition of spatial semantic categories. Language acquisition and conceptual development, 475-511.

Brier, G. W. (1950). Verification of forecasts expressed in terms of probability. Monthly Weather Review, 78(1), 1–3. https://doi.org/10.1175/1520-0493(1950)078<0001:VOFEIT>2.0.CO;2

Cannon, J., Levine, S., & Huttenlocher, J. (2007). A system for analyzing children and caregivers' language about space in structured and unstructured contexts. Spatial intelligence and learning center (SILC) technical report.

d'Apice, K., Latham, R. M., & von Stumm, S. (2019). A naturalistic home observational approach to children's language, cognition, and behavior. Developmental Psychology, 55(7), 1414.

Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019, June). BERT: Pre-training of deep bidirectional transformers for language understanding. In Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers) (pp. 4171-4186).

Fox, D. S., Ribner, A., Camarote, J., Vortuba-Drzal, E., Bachman, H. J., & Libertus, M. E. (2025). Parents' spatial talk predicts toddlers' spatial language gains. Learning and Instruction, 97, 102107. https://doi.org/10.1016/j.learninstruc.2025.102107

Frick, A. (2019). Spatial transformation abilities and their relation to later mathematics performance. Psychological research, 83(7), 1465-1484.

Guo, C., Pleiss, G., Sun, Y., & Weinberger, K. Q. (2017). On calibration of modern neural networks. In D. Precup & Y. W. Teh (Eds.), Proceedings of the 34th International Conference on Machine Learning (Vol. 70, pp. 1321–1330). PMLR. https://proceedings.mlr.press/v70/guo17a.html

Harris, D. (2023). Spatial reasoning in context: bridging cognitive and educational perspectives of spatial-mathematics relations. Frontiers in Education (Vol. 8, p. 1302099).

He, P., Liu, X., Gao, J., & Chen, W. (2021). DeBERTa: Decoding-enhanced BERT with disentangled attention. In International Conference on Learning Representations (ICLR 2021). https://openreview.net/forum?id=XPZIaotutsD

He, P., Gao, J., & Chen, W. (2023). DeBERTaV3: Improving DeBERTa using ELECTRA-style pre-training with gradient-disentangled embedding sharing. In The Eleventh International Conference on Learning Representations (ICLR 2023). https://openreview.net/forum?id=sE7-XhLxHA

Lauer, J. E., & Lourenco, S. F. (2016). Spatial processing in infancy predicts both spatial and mathematical aptitude in childhood. Psychological science, 27(10), 1291-1298.

Margolin, G., Oliver, P. H., Gordis, E. B., O'hearn, H. G., Medina, A. M., Ghosh, C. M., & Morland, L. (1998). The nuts and bolts of behavioral observation of marital and family interaction. Clinical child and family psychology review, 1(4), 195-213.

Noble, K. G., McCandliss, B. D., & Farah, M. J. (2007). Socioeconomic gradients predict individual differences in neurocognitive abilities. Developmental science, 10(4), 464-480.

Polinsky, N., Pagano, L. C., Acosta, D. I., Haden, C. A., & Uttal, D. H. (2023). Spatial language in families' conversational reflections about museum experiences. Journal of Applied Developmental Psychology, 86, 101539.

Pruden, S. M., Levine, S. C., & Huttenlocher, J. (2011). Children's spatial thinking: Does talk about the spatial world matter? Developmental Science, 14(6), 1417–1430. https://doi.org/10.1111/j.1467-7687.2011.01088.x

Raschka, S. (2018). Model evaluation, model selection, and algorithm selection in machine learning. arXiv preprint arXiv:1811.12808.

Shi, Q., & Uttal, D. (2025, October 27–29). Evaluating AI methods for coding spatial language from children's naturalistic conversations [Paper presentation]. Conference on Artificial Intelligence in Measurement and Education (AIME-Con), Pittsburgh, PA, United States.

Uttal, D. H., Meadow, N. G., Tipton, E., Hand, L. L., Alden, A. R., Warren, C., & Newcombe, N. S. (2013). The malleability of spatial skills: a meta-analysis of training studies. Psychological bulletin, 139(2), 352.

Yu, H., Yang, Z., Pelrine, K., Godbout, J. F., & Rabbany, R. (2023). Open, closed, or small language models for text classification?. arXiv preprint arXiv:2308.10092.

Zhang, Y., Jin, R., & Zhou, Z. H. (2010). Understanding bag-of-words model: a statistical framework. International journal of machine learning and cybernetics, 1(1), 43-52.

CDM Q-Matrix Discovery with LLMs: Fusing Domain Knowledge with Empirical Evidence

Susu Zhang
University of Illinois Urbana-Champaign
szhan105@illinois.edu

V.N. Vimal Rao
University of Illinois Urbana-Champaign
raovnv@illinois.edu

Abstract

We present a method for diagnostic assessment Q-matrix specification combining LLM input with model-based empirical validation. An LLM generates an initial Q-matrix from item content, and response data guide subsequent Q refinement. The approach integrates substantive rationale with empirical evidence to support scalable and measurement theory-grounded diagnostic assessment.

1 Introduction

Diagnostic assessments are designed to measure examinees' multidimensional profiles on fine-grained, discrete attributes, such as the skills, procedures, and knowledge required to solve assessment items (Rupp et al., 2010). Compared with a single continuous proficiency score, an attribute-mastery profile provides more specific information about students' strengths and weaknesses and may therefore support targeted remediation, instructional decisions, and personalized feedback. Cognitive diagnosis models (CDMs), also referred to as diagnostic classification models, provide a psychometric framework for relating these latent mastery profiles to observed item responses.

Central to diagnostic assessment is specifying the Q-matrix, a binary item-by-attribute incidence matrix indicating which attributes are required by each item. The Q-matrix specifies the relationship between attribute mastery and item responses and thereby determines the substantive interpretation of the resulting diagnostic classifications. Misspecification can affect both item parameter estimation and examinee classification (Rupp and Templin, 2008). A defensible Q-matrix should be grounded in both the cognitive demands implied by item content and empirical fit on the observed response data.

Traditional Q-matrix development relies on cognitive task analysis conducted by subject-matter specialists. These specialists map each item to a set of attributes and, where relevant, specify relationships such as attribute hierarchies. To support consistent coding by trained raters, the attributes are accompanied by clear labels, operational definitions, and rules for assigning item-attribute links. Think-aloud evidence from examinees can be used to assess the proposed attributes and to revise them when necessary. Support for a Q-matrix may therefore combine prior literature, expert ratings or panel discussion, evidence about examinees' solution processes, and subsequent validation using response data (Gierl et al., 2008; Li and Suen, 2013; Akbay et al., 2018). Coordinating these inputs can require subject-matter and measurement expertise, specialized data collection, and considerable time, making it costly for routine classroom assessments.

Empirically driven approaches aim to learn item-attribute relations from response data, addressing both the substantial human input required for expert construction and the possibility of misspecification despite careful review. These approaches are commonly divided into Q-matrix refinement methods and direct estimation methods (Liu and Kang, 2019; Zhang et al., 2023). Refinement methods begin with an expert-specified Q-matrix and use item-level diagnostics to identify potentially misspecified entries while preserving substantive interpretability. Examples include parameter-based validation under the deterministic-input, noisy-and-gate (DINA) model (Junker and Sijtsma, 2001; de la Torre, 2008), residual-based refinement (Chiu, 2013), and extensions to more general CDMs or methods that favor parsimonious revisions through fit-parsimony criteria or regularization (de la Torre and Chiu, 2016; Nájera et al., 2021). Expert judgment is therefore involved both in constructing the starting matrix and in deciding whether proposed revisions are substantively defensible. Direct estimation methods infer an unknown Q-matrix from response data (e.g., Liu et al., 2012, 2013; Chen et al., 2018). Dependable recovery may require

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 216–230
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

more sample information than is available in small or moderately sized classroom datasets, especially early in an assessment's deployment. The estimated binary incidence matrix also requires post-estimation interpretation to identify the cognitive processes represented by its columns and explain particular item-attribute associations: Similar to exploratory factor analysis, the columns (here, attributes) are not inherently interpretable.

Recent work has examined the emerging capabilities of general-purpose large language models (LLMs) for psychometric tasks, including automated scoring, automated item generation, and response simulation (e.g., Gorgun and Bulut, 2024; Suk et al., 2025). Conceptual and review papers have also discussed their potential utility for formative assessment, including assistance with item writing and more timely, customized instructional feedback (Chang, 2024; Hao et al., 2024; Ho, 2024). Meanwhile, LLM outputs vary across repeated generations and prompt specifications, and their accuracy requires empirical verification when they inform measurement decisions.

Within cognitive diagnosis, recent studies have examined content-based Q-matrix construction with LLMs from item and attribute descriptions and evaluated the resulting matrices through expert comparison and empirical model fit (Aşiret and Sünbül, 2025; Xue and Appleton, 2026). Two key challenges remain for classroom use. First, an LLM's content-based proposal needs to be evaluated against examinee response patterns and revised when the proposed item-attribute relations are inconsistent with those patterns. Second, the initial specification and each empirical revision need to remain interpretable so that teachers and domain experts can evaluate whether the proposed relations reflect the skills required by the items.

The current work draws from automated statistical model discovery with language models (Li et al., 2024; Zhang et al., 2025) and proposes a method for Q-matrix discovery under the DINA model that integrates LLM-generated proposals based on domain knowledge and item content with empirical evidence from examinee responses. The LLM first proposes a Q-matrix from the item content and attribute descriptions and provides a brief rationale for each item Q-vector. After the DINA model is fitted, an item-level Q-matrix validation statistic identifies alternative Q-vectors that are more consistent with the observed responses. These alternatives are returned to the LLM for evaluation against the item content, and the LLM provides a rationale for each decision to revise or retain the flagged Q-vector. Global model fit is then used to select among the Q-matrices visited during refinement. This study makes the following contributions:

- An LLM-assisted algorithm for initializing the Q-matrix from item content and refining the initial specification using empirical response evidence under the DINA model;
- Support for substantive interpretation and expert review through brief LLM-generated rationales for the initial item Q-vectors and subsequent refinement decisions;
- Evaluation of the procedure with (1) a proof-of-concept simulation study and (2) an empirical application from a large university general-education statistics course. The empirical evaluation combines quantitative evaluation (e.g., expert agreement, model fit) with expert evaluation of the LLM-generated rationales.

2 Method

Let $\mathbf{Y} = (Y_{ij})$ be the $N \times J$ matrix of binary responses and $\mathbf{Q} = (q_{jk})$ the $J \times K$ binary Q-matrix, where $q_{jk} = 1$ indicates that item j requires attribute k, and $q_{jk} = 0$ otherwise. The procedure begins with an LLM proposal based on item content and then alternates between Bayesian DINA estimation, item-level Q-vector validation, and LLM refinement of the Q-vector on empirically flagged items.

2.1 Initial Q-Matrix Proposal

To obtain the initial LLM-proposed Q-matrix, we designed a prompt around three components recommended for LLM-based Q-matrix construction: a clearly specified role and task, the content needed to make item-attribute judgments, and explicit constraints on the requested output (Xue and Appleton, 2026). The LLM was assigned the role of a subject-matter expert for the assessment domain and target population and asked to construct an initial Q-matrix. For each item, it received the stem, response options, correct answer, and any table or figure needed to solve the item. Each attribute was represented by its label, definition, and examples in the assessment domain.

The prompt's item-coding instructions were informed by expert Q-matrix development practices, in which an item is solved, the necessary steps are

identified, and those steps are mapped to the specified attributes (Gierl et al., 2008; Rupp et al., 2010). The LLM was therefore asked to infer a defensible solution and code only attributes required by that solution. The prompt provided brief coding guidance, distinguishing required skills from incidental/helpful skills, and imposing additional constraints (e.g., attribute coverage, number of single-attribute items, and distinguishable columns) for DINA model identifiability (Gu and Xu, 2019). The full prompt and application-specific constraints are reported in Appendix F.

The LLM returned a binary $J \times K$ Q-matrix and a brief rationale for every item's Q-vector. Each rationale linked the coded attributes to the steps that made them necessary, and a final consistency check required the rationale and Q-vector to identify the same attributes. Although only the Q-matrix is used for subsequent model fit, the rationale provides an explanation to the to the initial proposal and allowed substantive review.

2.2 Bayesian DINA Estimation and Q-Validation

Given a proposed Q-matrix, a DINA model is fitted using the Bayesian Gibbs sampling algorithm of Culpepper (2015). The DINA model is a simple conjunctive model, where an examinee answers an item correctly with high probability only if they mastered all requisite skills of the item. For examinee i and item j, let $\boldsymbol{\alpha}_i$ denote the binary attribute profile and define the ideal response

$$\eta_{ij}(\boldsymbol{\alpha}_i, \mathbf{q}_j) = \prod_{k=1}^{K} \alpha_{ik}^{q_{jk}}. \tag{1}$$

The conditional correct-response probability is

$$P(Y_{ij} = 1 \mid \boldsymbol{\alpha}_i, \mathbf{q}_j) = (1 - s_j)^{\eta_{ij}} g_j^{1-\eta_{ij}}, \tag{2}$$

where s_j and g_j are the slipping and guessing parameters. Each chain was run for 10,000 iterations with the first 5,000 as burn-in.

In practice, item-level Q-matrix validation statistics identify items for expert review and suggest alternative Q-vectors that agree more closely with the response data. Experts can then examine whether the suggested attributes are supported by the item content and revise the Q-matrix accordingly. We implemented this process using a Bayesian adaptation of the residual sum of squares (RSS) statistic proposed for the DINA model by Chiu (2013) for Q-matrix validation. For every item j and every nonzero candidate Q-vector $\mathbf{q}_m \in \{0,1\}^K$, let $\eta_{im}^{(s)}$ denote the ideal response implied by the sampled attribute profile $\boldsymbol{\alpha}_i^{(s)}$ at posterior draw s. We calculated

$$\widehat{\mathrm{RSS}}_{jm} = \frac{1}{S} \sum_{s=1}^{S} \sum_{i=1}^{N} \left(Y_{ij} - \eta_{im}^{(s)} \right)^2. \tag{3}$$

Intuitively, $\widehat{\mathrm{RSS}}_{jm}$ captures the posterior discrepancy between the ideal responses implied by candidate Q-vector $\mathbf{q}_m$ and the observed responses on item j, where smaller values indicate closer agreement with the observed item responses.

Relative overall fit was evaluated using the widely applicable information criterion (WAIC; Watanabe, 2010). For posterior draws $\boldsymbol{\Omega}^{(1)}, \ldots, \boldsymbol{\Omega}^{(S)}$ and examinee response vector $\mathbf{y}_i$,

$$\mathrm{WAIC} = -2\left(\mathrm{lppd} - p_{\mathrm{WAIC}}\right), \text{ where} \tag{4}$$

$$\mathrm{lppd} = \sum_{i=1}^{N} \log \left[\frac{1}{S} \sum_{s=1}^{S} p(\mathbf{y}_i \mid \boldsymbol{\Omega}^{(s)}) \right], \tag{5}$$

$$p_{\mathrm{WAIC}} = \sum_{i=1}^{N} \mathrm{Var}_s \left[\log p(\mathbf{y}_i \mid \boldsymbol{\Omega}^{(s)}) \right]. \tag{6}$$

Lower WAIC values indicate better expected predictive fit relative to the other fitted Q-matrices.

2.3 Model Critique and Refinement

After each fit, a critique was generated for the previously visited Q-matrix with the lowest WAIC. An item was flagged whenever at least one candidate Q-vector had a lower RSS than its current vector. For each flagged item, the current vector, the RSS-minimizing alternative, and their differences in RSS values were presented as an empirical response-based critique of the current Q-vector.

The refinement prompt included only the flagged items, accompanied by their item content, current Q-vectors, attribute definitions, structural constraints, and RSS-based statistical critique. The LLM was instructed to solve each flagged item, decide which attributes were required, compare that with the RSS-suggested vectors, and avoid accepting a change solely because it reduced RSS. The LLM could retain the current Q-vector, adopt the RSS-minimizing vector, or return a different content-based or partially-revised vector when only part of the RSS-suggested change was supported by the item content.

In addition to the revised Q, for each flagged item, the refinement prompt instructed the LLM to provide a short rationale that began by indicating whether the current Q-vector was changed or unchanged and then explained that decision, i.e., the item content or solution step supporting the revised attribute assignment, or why the current coding better represented the item content than the RSS-suggested alternative. The prompt also instructed the LLM to perform a self-consistency check so that the rationale was consistent with attributes required by the revised Q.

To ensure that the search advanced, each refinement call was required to alter at least one flagged row. The complete refinement prompt template is provided in Appendix G. The refined Q-matrix is subsequently fit to the response data to obtain the RSS-based critique for the next iteration.

2.4 Stopping and Final Selection

At each iteration, the critique and refinement were based on the lowest-WAIC Q-matrix visited so far. The maximum number of iterations was set as 10, and the algorithm stops earlier if no item was flagged or the same Q-matrix was produced in three consecutive iterations. The matrix with the lowest WAIC across all visited iterations was selected as the final refined Q-matrix.

2.5 LLM Implementation Settings

All LLM calls used GPT-5.4 through API call with medium reasoning effort and temperature 1. GPT-5.4 was selected because it supported both textual and image inputs and inference-time reasoning. Comparison of foundation models was outside the scope of the study. The same model settings were used for initial proposals and refinements in both applications. API responses were requested as R objects, checked for formatting, and regenerated after a parsing or output-validation failure.

3 Fraction-Subtraction Simulation

3.1 Simulation Design

We examine the proposed methods with a proof-of-concept simulation study, using the item content of the 20 fraction-subtraction items originally reported by Tatsuoka (1984), and a four-attribute ground-truth Q-matrix for the fraction items reported by Chen et al. (2018), who learned the matrix from empirical response data. Köhn et al. (2025) subsequently interpreted the four attributes as finding a common denominator, borrowing from an integer, subtracting integer and fractional parts, and converting a whole number to a fraction. Among the Qs proposed for the fraction-subtraction items, we used this specification because it met the DINA identifiability conditions.

We conducted 50 simulation replications. In each replication, we generated a binary response matrix for 536 examinees and 20 items under the DINA model using the four-attribute ground-truth Q-matrix. The sample size matched the empirical fraction-subtraction dataset, item-specific slipping and guessing parameters were set to the estimates reported by Köhn et al. (2025), and true attribute profiles were sampled uniformly from the 2^4 classes. In each replication, the LLM received the original item content and answer key together with the four attribute definitions and examples to propose an initial Q-matrix. Subsequent refinement incorporated the response data generated under the ground-truth Q. The inputs supplied to the LLM are reported in Appendix A. Note that the item text was developed independently of the data-generating Q-matrix, and the item content itself does not have a known ground-truth Q-matrix. The evaluation therefore examined whether incorporating response evidence enabled the algorithm to recover the Q-matrix used to generate the simulated responses.

3.2 Evaluation Criteria

For each replication, Q-matrix recovery was evaluated for the initial LLM proposal and the final matrix selected by WAIC. Entry-wise accuracy was the proportion of all $J \times K$ entries correctly recovered. Sensitivity was the proportion of true entries equal to 1 that were recovered as 1, and specificity was the proportion of true entries equal to 0 that were recovered as 0. We report the mean and standard deviation of each measure across replications. For each item, we also calculated the proportion of replications in which it was flagged by RSS, the proportion of flags followed by an LLM revision, and the proportion of replications in which the final estimated Q-vector matched the data-generating Q-vector.

3.3 Results

As shown in Figure 1, global model fit and Q-matrix recovery generally improved over refinement iterations. Table 1 shows that mean entry-wise accuracy increased from .869 for the initial proposals to .975 for the final matrices. Mean sen-

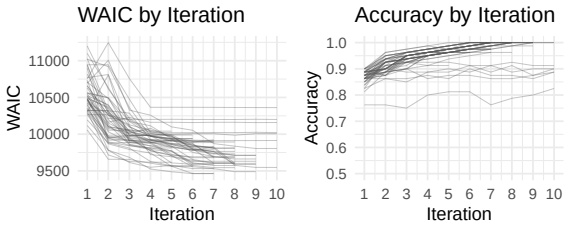


Figure 1: Fraction simulation under DINA: WAIC and Q-matrix entry-wise accuracy across refinement iterations. Each line represents a replication.

	Accuracy	Sensitivity	Specificity
Initial Q	.869	.785	.922
SD	(.026)	(.052)	(.020)
Final Q	.975	.963	.983
SD	(.040)	(.057)	(.032)

Table 1: Recovery of the data-generating Q-matrix across 50 fraction-subtraction simulation replications.

sitivity increased from .785 to .963, and specificity increased from .922 to .983. Item-level results are reported in Appendix B. Although every replication used the same item content, attribute definitions, and LLM settings, many entries of the initial Q-matrices varied across replications, producing mean entry values farther from 0 or 1. The final Q-matrices contained fewer entries with mean absolute deviations greater than .20 from the data-generating values, shown in bold in Appendix B, indicating higher item-level accuracy. Their mean entry values were also more concentrated near 0 or 1, indicating lower variation across replications after incorporating response evidence.

4 Statistics Classroom Assessment

4.1 Data Description

The empirical application analyzed the first midterm examination administered in Fall 2025 in a large general-education introductory statistics course at a Midwestern public university. The examination occurred in Week 5 of a 15-week term and assessed concepts involving foundational principles of data, measurement, and descriptive statistics. Its 25 multiple-choice items covered data structure, variable types, relative and conditional frequencies, numerical summaries and visualizations, characteristics of quantitative distributions, comparisons across groups, and associations represented in two-way tables and scatterplots.

The course instructor identified eight attributes that the examination was intended to measure and specified two alternative organizations of those attributes. The attribute labels and definitions are detailed in Table 8. In version 1, A3 represented summaries and graphs for a single quantitative variable, whereas A4 represented categorical frequencies and bar graphs. In version 2, A3 represented numerical summaries across quantitative and categorical variables, whereas A4 represented visual displays across both variable types. The definitions of A1, A2, and A5–A8 were identical. The instructor also specified a corresponding expert Q-matrix under each version, serving as a substantive reference for comparison with the learned Q-matrices. The items were developed from a blueprint aligned with the course content and learning objectives and had been refined across semesters using item response theory analyses of student response data. Appendix C provides item summaries while protecting content that remains in operational use.

For each attribute version, the LLM received the same 25 item stems, response options, and correct answers, together with any table or plot required to solve an item and the corresponding version of the eight attribute definitions. The empirical data contained binary responses from 1,144 students. The same 1144×25 response matrix was used to fit and evaluate the candidate Q-matrices under the DINA model for both attribute versions.

4.2 Evaluation Criteria

The initial and selected final Q-matrices under each attribute version were compared with the corresponding expert Q-matrix. Treating expert-coded entries as the reference, we calculated entry-wise agreement, sensitivity, and specificity. We also fitted the expert Q-matrices using the same Bayesian DINA procedure and compared their WAIC values with those of the LLM initial and final Q-matrices.

For the best-fitting Q, two experts who are university-level statistics instructors independently rated the rationales using four binary criteria that assessed whether each rationale matched the attributes coded in the Q-vector (*attribute alignment*), was factually grounded in the item content (*item grounding*), agreed with the rater's judgment of the attributes required or not required to solve the item (*substantive accuracy*), and was sufficiently clear and specific for a reviewer to understand the initial coding or refinement decision (*review utility*). Detailed rating instructions and examples of violations are provided in Appendix D.

The initial Q rationale evaluation included all 25

items. The refinement rationale evaluation included the last available rationale for each item flagged along the trajectory leading to the final selected Q-matrix. For each criterion, endorsement was calculated as the proportion of ratings equal to 1 across both raters; inter-rater agreement was the proportion of items receiving identical ratings for the two raters.

4.3 Results

Version 1 terminated at 10 iterations and selected iteration 6 by WAIC, where 10 items were flagged in at least one iteration, and 7 items' Q-vectors differed between the initial and final Q-matrices. Version 2 terminated after five iterations and selected iteration 2, and six items were flagged and three item Q-vectors differed. As shown in Table 2, the final refined Q-matrix had a lower WAIC than both the initial LLM Q-matrix and the expert Q-matrix for both attribute versions. The empirically refined Q under version 1 had the lowest WAIC. The final matrices also had slightly greater agreement with their corresponding expert matrices than the initial proposals. Note that whereas specificity increased, sensitivity was unchanged under version 2 and decreased under version 1.

Measure	Q-matrix	Version 1	Version 2
WAIC	Initial LLM	29167.21	29145.52
	Final refined	**29141.98**	**29142.41**
	Expert	29148.19	29145.87
Agreement	Initial LLM	.920	.915
	Final refined	**.925**	**.925**
Sensitivity	Initial LLM	**.815**	**.741**
	Final refined	.741	**.741**
Specificity	Initial LLM	.936	.942
	Final refined	**.954**	**.954**

Table 2: Bayesian DINA fit and agreement with the expert Q-matrix for the statistics assessment.

The LLM initial and final refinement rationales under version 1 are reported in Appendix E and were evaluated by both experts on the four criteria. Table 3 reports the average endorsement and agreement of the expert ratings. The rationales were consistently rated to be grounded in the item content. Initial-rationale endorsement ranged from .92 to 1.00 across criteria, with inter-rater agreement from .92 to 1.00. Among the 10 refinement rationales, endorsement ranged from .90 to 1.00 and agreement from .80 to 1.00.

	Initial ($J = 25$)		Refinement ($J = 10$)	
Criterion	Endorse.	Agree.	Endorse.	Agree.
Alignment	1.00	1.00	.95	.90
Grounding	1.00	1.00	1.00	1.00
Accuracy	.92	.92	.95	.90
Review utility	.98	.96	.90	.80

Table 3: Expert evaluation of version 1 LLM-generated rationales. Endorsement: proportion of binary ratings equal to 1 across two experts; agreement: agreement rate between raters across items.

5 Discussion

The current study proposes an iterative approach to Q-matrix discovery that integrates LLM judgments based on item content with empirical evidence from item responses. The LLM proposes an initial Q-matrix from the items and attribute definitions, and diagnostics from the fitted Bayesian DINA model are used to generate statistical critique and inform subsequent revisions. Additionally, to maintain interpretability of the learned Qs, the LLM provides brief, item-specific rationales for both the initial attribute assignments and later decisions to revise or retain them. These rationales allow instructors and domain experts to examine the basis for each Q-vector, assess whether the coded attributes reflect the intended item requirements, and evaluate discrepancies between the content-based proposal and empirical diagnostic evidence.

Results from the fraction-subtraction simulation indicate that empirical refinement substantially improved recovery of the data-generating Q-matrix relative to the initial LLM specifications. Each replication used the same item and attribute descriptions, yet stochastic variation in the LLM outputs produced different initial Q-matrices. Incorporating response-based diagnostics led to more consistent final specifications, suggesting that empirical evidence may stabilize content-based proposals in addition to improving average recovery.

For the statistics classroom assessment, the final refined Q-matrix had lower WAIC than the initial LLM Q-matrix and the expert Q-matrix under both attribute versions, indicating better expected predictive fit among the fitted Q-matrices. Refinement also produced a small increase in agreement with the instructor-specified Q-matrices, associated primarily with higher specificity. The expert ratings of the rationales indicated high overall alignment, grounding, substantive accuracy, and review utility for both sets of rationales. For items on which the

expert and refined Q-vectors differ, the item-level rationales and fit diagnostics provide information for evaluating whether the discrepancy reflects misspecification, an alternative solution strategy, or a difference in attribute granularity.

There are several limitations of this study. First, the procedure holds the prespecified attribute set fixed, leaving possible missing, redundant, or overly broad/narrow attributes unexamined. Future extensions could use recurring conflicts between content-based judgments and response-based diagnostics, sparse attribute coverage, or weak empirical item discrimination to prompt structured review of the attribute definitions and proposals to merge, split, remove, redefine, or add attributes, followed by expert review. Second, refinement is conducted under the conjunctive DINA model, so suggested Q-matrix revisions may also reflect misspecification of how attributes interact to affect correct response probability. Future research can extend the workflow to more general CDMs. Finally, the classroom results were based on a single assessment and examinee sample of a single course, and only one foundation model was examined (GPT-5.4). Generalizability should be evaluated across domains, item formats, sample sizes, prompts, and LLMs.

Additional extensions could incorporate attribute hierarchies, with admissible mastery profiles and Q-matrix revisions constrained by prerequisite relations among attributes. The method could also be adapted to open-ended responses, allowing response content, written work, or solution traces to supplement correctness scores when evaluating the proposed item-attribute relations. Such evidence could help experts assess whether the attributed skills correspond to the processes examinees actually used.

As LLM capabilities continue to evolve, the current study provides a practical workflow to harness LLM capabilities to support instructors in implementating of measurement-theory-based best practices in diagnostic classroom assessment.

References

Lokman Akbay, Ragip Terzi, Mehmet Kaplan, and Katibe Gizem Karaaslan. 2018. Expert-based attribute identification and validation on fraction subtraction: A cognitively diagnostic assessment application. *Journal on Mathematics Education*, 9(1):103–120.

Semih Aşiret and Seçil Ömür Sünbül. 2025. Evaluating ChatGPT-4o's performance in construction of Q-matrix for a cognitive diagnostic assessment. *Research Square*.

Hua-Hua Chang. 2024. Harnessing AI for educational measurement: Standards and emerging frontiers. *J. Educ. Behav. Stat.*, 49(5):702–708.

Yinghan Chen, Steven Andrew Culpepper, Yuguo Chen, and Jeffrey Douglas. 2018. Bayesian estimation of the DINA Q-matrix. *Psychometrika*, 83(1):89–108.

Chia-Yi Chiu. 2013. Statistical refinement of the Q-matrix in cognitive diagnosis. *Applied Psychological Measurement*, 37(8):598–618.

Steven Andrew Culpepper. 2015. Bayesian estimation of the dina model with gibbs sampling. *Journal of Educational and Behavioral Statistics*, 40(5):454–476.

Jimmy de la Torre. 2008. An empirically based method of Q-matrix validation for the DINA model: Development and applications. *J. Educ. Meas.*, 45(4):343–362.

Jimmy de la Torre and Chia-Yi Chiu. 2016. A general method of empirical Q-matrix validation. *Psychometrika*, 81(2):253–273.

Mark J Gierl, Changjiang Wang, and Jiawen Zhou. 2008. Using the attribute hierarchy method to make diagnostic inferences about examinees' cognitive skills in algebra on the sat. *Journal of Technology, Learning, and Assessment*, 6(6):n6.

Guher Gorgun and Okan Bulut. 2024. Instruction-tuned large-language models for quality control in automatic item generation: A feasibility study. *Educ. Meas. Issu. Pr.*

Yuqi Gu and Gongjun Xu. 2019. The sufficient and necessary condition for the identifiability and estimability of the dina model. *Psychometrika*, 84(2):468–483.

Jiangang Hao, Alina A von Davier, Victoria Yaneva, Susan Lottridge, Matthias von Davier, and Deborah J Harris. 2024. Transforming assessment: The impacts and implications of large language models and generative AI. *Educ. Meas. Issu. Pr.*, 43(2):16–29.

Andrew D Ho. 2024. Artificial intelligence and educational measurement: Opportunities and threats. *J. Educ. Behav. Stat.*, 49(5):715–722.

Brian W Junker and Klaas Sijtsma. 2001. Cognitive assessment models with few assumptions, and connections with nonparametric item response theory. *Applied Psychological Measurement*, 25(3):258–272.

Hans-Friedrich Köhn, Chia-Yi Chiu, Olasumbo Oluwalana, Hyunjoo Kim, and Jiaxi Wang. 2025. A two-step Q-matrix estimation method. *Applied Psychological Measurement*, 49(1-2):3–28.

Hongli Li and Hoi K Suen. 2013. Constructing and validating a Q-matrix for cognitive diagnostic analyses of a reading test. *Educ. Assess.*, 18(1):1–25.

Michael Y Li, Emily B Fox, and Noah D Goodman. 2024. Automated statistical model discovery with language models. In *Proceedings of the 41st International Conference on Machine Learning*, pages 27791–27807.

Jingchen Liu and Hyeon-Ah Kang. 2019. Q-matrix learning via latent variable selection and identifiability. In *Handbook of Diagnostic Classification Models*, pages 247–263. Springer.

Jingchen Liu, Gongjun Xu, and Zhiliang Ying. 2012. Data-driven learning of Q-matrix. *Applied Psychological Measurement*, 36(7):548–564.

Jingchen Liu, Gongjun Xu, and Zhiliang Ying. 2013. Theory of the self-learning Q-matrix. *Bernoulli*, 19(5A):1790–1817.

Pablo Nájera, Miguel A Sorrel, Jimmy de la Torre, and Francisco José Abad. 2021. Balancing fit and parsimony to improve Q-matrix validation. *Br. J. Math. Stat. Psychol.*, 74(S1):110–130.

Andr'e A. Rupp and Jonathan Templin. 2008. The effects of q-matrix misspecification on parameter estimates and classification accuracy in the dina model. *Educational and Psychological Measurement*, 68(1):78–96.

Andre A. Rupp, Jonathan Templin, and Robert A. Henson. 2010. *Diagnostic measurement: Theory, methods, and applications*. The Guilford Press, New York, NY.

Youmi Suk, Chenguang Pan, and Ke Yang. 2025. Using generative AI for sequential data generation in monte carlo simulation studies. *J. Educ. Behav. Stat.*, (10769986251397559).

Kikumi K. Tatsuoka. 1984. *Analysis of errors in fraction addition and subtraction problems*. Computer-Based Education Research Laboratory, University of Illinois at Urbana-Champaign, Urbana, IL. Final Report #143; Final report for grant no. NIE-G-81-0002.

Sumio Watanabe. 2010. Asymptotic equivalence of Bayes cross validation and widely applicable information criterion in singular learning theory. *Journal of Machine Learning Research*, 11(12).

Kang Xue and James J Appleton. 2026. Evaluating general-purpose multimodal AI for q-matrix generation from math items: A cognitive diagnostic modeling exploration. *J. Educ. Meas.*, 63(1):e70028.

Erica Zhang, Ryunosuke Goto, Naomi Sagan, Jurik Mutter, Nick Phillips, Ash Alizadeh, Kangwook Lee, Jose Blanchet, Mert Pilanci, and Robert Tibshirani. 2025. Llm-lasso: A robust framework for domain-informed feature selection and regularization. *arXiv preprint arXiv:2502.10648*.

Susu Zhang, Jingchen Liu, and Zhiliang Ying. 2023. Statistical applications to cognitive diagnostic testing. *Annual Review of Statistics and Its Application*, 10:651–675.

A Fraction-Subtraction Items and Attributes

Table 4 displays the item texts and answer keys supplied to the LLM. Table 5 shows the four attributes' definitions and examples.

Item	Expression	Accepted answer
1	$5/3 - 3/4$	$11/12$
2	$3/4 - 3/8$	$3/8$
3	$5/6 - 1/9$	$13/18$
4	$3\ 1/2 - 2\ 3/2$	0
5	$4\ 3/5 - 3\ 4/10$	$6/5$ or $1\ 1/5$
6	$6/7 - 4/7$	$2/7$
7	$3 - 2\ 1/5$	$4/5$
8	$2/3 - 2/3$	0
9	$3\ 7/8 - 2$	$15/8$ or $1\ 7/8$
10	$4\ 4/12 - 2\ 7/12$	$7/4$ or $1\ 3/4$
11	$4\ 1/3 - 2\ 4/3$	1
12	$1\ 1/8 - 1/8$	1
13	$3\ 3/8 - 2\ 5/6$	$13/24$
14	$3\ 4/5 - 3\ 2/5$	$2/5$
15	$2 - 1/3$	$5/3$ or $1\ 2/3$
16	$4\ 5/7 - 1\ 4/7$	$22/7$ or $3\ 1/7$
17	$7\ 3/5 - 4/5$	$34/5$ or $6\ 4/5$
18	$4\ 1/10 - 2\ 8/10$	$13/10$ or $1\ 3/10$
19	$4 - 1\ 4/3$	$5/3$ or $1\ 2/3$
20	$4\ 1/3 - 1\ 5/3$	$5/3$ or $1\ 2/3$

Table 4: Fraction-subtraction item information supplied to the LLM.

Attribute	Definition and example
A1: Find a common denominator	Construct a common denominator and rewrite both terms as equivalent fractions with that denominator. Example: $5/6$ and $1/9$ become $15/18$ and $2/18$.
A2: Borrow from an integer	When the fractional part of the minuend is smaller, borrow 1 from its integer part and add d/d to its fractional part. Example: $4\ 1/10$ becomes $3\ 11/10$.
A3: Subtract integer and fractional parts	Subtract the integer and fractional parts separately after placing them in compatible form, then combine the result. Example: $3\ 16/12 - 2\ 7/12 = 1\ 9/12$.
A4: Convert a whole number to a fraction	Rewrite a whole-number operand as a fraction with the working denominator. Example: rewrite 2 as $6/3$ when computing $2 - 1/3$.

Table 5: Fraction-subtraction attribute definitions supplied to the LLM.

B Fraction-Subtraction Simulation Item-Level Results

	Initial Q				Final Q						
	A1	A2	A3	A4	A1	A2	A3	A4	Row-wise Q Recovery	RSS-based flag rate	Revision rate given RSS flag
Item1: 5/3 - 3/4	1.000	0.000	0.000	0.000	1.000	0.000	0.000	0.000	1.000	0.120	0.000
Item2: 3/4 - 3/8	1.000	0.000	0.000	0.000	1.000	0.000	0.000	0.000	1.000	0.080	0.000
Item3: 5/6 - 1/9	1.000	0.000	0.000	0.000	1.000	0.000	0.000	0.000	1.000	0.120	0.000
Item4: 3 1/2 - 2 3/2	0.000	0.920	**0.220**	0.020	0.000	1.000	0.100	0.000	0.900	0.200	0.824
Item5: 4 3/5 - 3 4/10	1.000	0.000	1.000	0.000	1.000	0.000	1.000	0.000	1.000	0.020	0.000
Item6: 6/7 - 4/7	**1.000**	0.000	**0.000**	0.000	0.000	0.000	1.000	0.000	1.000	1.000	0.758
Item7: 3 - 2 1/5	0.000	0.140	**0.280**	0.900	0.000	0.100	0.120	0.940	0.880	0.400	0.554
Item8: 2/3 - 2/3	**1.000**	0.000	**0.000**	0.000	0.000	0.000	1.000	0.000	1.000	1.000	0.769
Item9: 3 7/8 - 2	0.000	0.000	1.000	0.000	0.000	0.000	1.000	0.000	1.000	0.000	-
Item10: 4 4/12 - 2 7/12	0.000	1.000	**1.000**	0.000	0.000	1.000	**0.340**	0.000	0.660	0.840	0.230
Item11: 4 1/3 - 2 4/3	0.000	0.880	**0.600**	0.020	0.000	0.980	**0.760**	0.020	0.740	0.260	0.944
Item12: 1 1/8 - 1/8	0.000	0.000	1.000	0.000	0.000	0.000	1.000	0.000	1.000	0.000	-
Item13: 3 3/8 - 2 5/6	1.000	1.000	0.980	**0.000**	1.000	1.000	1.000	0.840	0.840	1.000	0.248
Item14: 3 4/5 - 3 2/5	0.000	0.000	1.000	0.000	0.000	0.000	1.000	0.000	1.000	0.000	-
Item15: 2 - 1/3	0.000	0.000	0.000	1.000	0.000	0.020	0.020	0.980	0.980	0.120	0.025
Item16: 4 5/7 - 1 4/7	0.000	0.000	1.000	0.000	0.000	0.000	1.000	0.000	1.000	0.000	-
Item17: 7 3/5 - 4/5	0.000	1.000	**0.460**	0.000	0.000	1.000	**0.720**	0.000	0.720	0.260	0.933
Item18: 4 1/10 - 2 8/10	0.000	1.000	1.000	**0.000**	0.000	1.000	1.000	0.840	0.840	1.000	0.196
Item19: 4 - 1 4/3	0.000	**0.340**	**0.480**	0.940	0.000	0.980	0.960	0.940	0.900	0.700	0.722
Item20: 4 1/3 - 1 5/3	0.000	0.840	0.980	0.120	0.000	0.900	1.000	0.100	0.900	0.080	1.000

Table 6: Fraction simulation: Average initial Q across replications, average final Q across replications and final Q's row-wise recovery rate, proportion of replications an item is flagged at least once with RSS (i.e., RSS-based flag rate), and proportion of times Q vector is revised following RSS flag (i.e., Revision rate given RSS flag). *Note:* **Bolded** entries indicate an absolute deviation from ground-truth Q above .20, which could arise with either fluctuation across replications or consistent difference from ground-truth.

C Statistics Classroom Assessment Items and Attributes

Table 7 summarizes the operational items without reproducing the exact content of the item. Table 8 presents the attribute labels and definitions supplied to the LLM for Version 1 and Version 2. A1, A2, and A5–A8 were common to both versions, and A3 and A4 were version-specific.

Item	Description	Reference	p(correct)
1	Definition of data	-	.830
2	Number of attributes in a data set about hotels	Table	.943
3	Number of observations in a data set about hotels	Table	.950
4	Graph choice for hotel floor counts	Table	.755
5	Definition of relative frequency	-	.810
6	Computation of relative frequency	Table	.751
7	Quartile interpretation from a summary table	Table	.649
8	Identification of skewness from a summary table	Table	.739
9	Interpretation of a percentile from a summary table	Table	.757
10	Interpretation of the middle 50% from a summary table	Table	.676
11	Interpretation of the middle 95% from a summary table	Table	.495
12	Graph choice for in-state percentage by college	Table	.596
13	Denominator for a conditional frequency from a two-way table	Table	.506
14	Interpretation of Cramer's V from a two-way table	Table	.686
15	Definition of a smoothed trend line	Plot	.441
16	Interpretation of association strength in a scatterplot	Plot	.778
17	Interpretation of association direction in a scatterplot	Plot	.217
18	Identification of the variable types represented in a scatterplot	-	.751
19	Graph choice for comparing a quantitative distribution between groups	Plot	.898
20	Comparison of spread between two quantitative distributions	Plot	.960
21	Interpretation of an effect size	Plot	.413
22	Comparison of spread between two quantitative distributions	Plot	.890
23	Interpretation of distribution shape	Plot	.866
24	Interpretation of the center of a distribution	Plot	.537
25	Interpretation of the middle 95% of a distribution	Plot	.698

Table 7: Item summary for the statistics classroom assessment. *Note:* Reference: additional table or plot applicable to item. p(correct): Observed proportion of correct response.

Version	Attribute	Label	Definition
Both	A1	Data Structure	Systematic arrangement of a data set in terms of data points (observations) and their characteristics (attributes or variables).
Both	A2	Variable Types	Types of variables and measures, including quantitative and categorical variables and common subtypes.
1	A3	Univariate Quantitative Summaries and Graphs	Univariate summary statistics and common visual displays for a single quantitative variable, including the mean, median, standard deviation, five-number summary, density plot, and box plot.
1	A4	Univariate Categorical Frequencies and Bar Graphs	Frequency tables and visual displays for a single categorical variable, including frequencies, relative frequencies, percents, and bar graphs.
2	A3	Univariate Summaries	Univariate numerical summaries for quantitative and categorical variables, including the mean, median, standard deviation, five-number summary, frequencies, relative frequencies, and percents.
2	A4	Univariate Visualizations	Univariate visual displays for quantitative and categorical variables, including density plots, box plots, and bar graphs.
Both	A5	Quantitative Distribution Characteristics	Characteristics of quantitative distributions, including shape, location, spread, skewness, central intervals, and the normal rule of thumb.
Both	A6	Two-Way Tables	Two-way tables for two categorical variables, including conditional percents, comparisons of conditional percents, and association summaries such as Cramer's V.
Both	A7	Comparing Quantitative Distributions	Comparisons of quantitative distributions across groups, including differences in location and spread and standardized comparisons such as effect size and ratios of standard deviations.
Both	A8	Scatterplots and Correlations	Joint distributions and associations for two quantitative variables, including scatterplots, correlation-type measures, and trend lines.

Table 8: Statistics classroom assessment attribute labels and definitions supplied to the LLM.

D Rationale Rating Rubric

Criterion	Definition	Example of a violation
Attribute alignment	The rationale connects each explanation to the correct named attribute and does not confuse one attribute with another.	The rationale assigns a solution step to an attribute whose definition describes a different skill.
Item grounding	The rationale remains tied to the item content and solution steps and does not introduce unsupported details.	For an item that only requires comparing fractions, the rationale introduces borrowing or mixed-number conversion even though neither operation appears in the item.
Substantive accuracy	The rationale correctly explains whether each attribute is or is not required to solve the item.	The rationale identifies an attribute as required even though the item can be solved without it, or rejects an attribute that is necessary.
Review utility	The rationale is clear and specific enough for a reviewer to understand the initial Q-vector or the decision to revise or retain it.	The rationale states only that a skill is relevant without identifying what feature or solution step makes it necessary or unnecessary.

Table 9: Instructions used to evaluate the LLM-generated rationales.

E Version 1 LLM Rationales Evaluated by Experts

Table 10 contains the LLM's rationale for the initial Q-vector for all 25 items. Table 11 contains the last available refinement rationale for each of the 10 items flagged on the trajectory leading to the selected version 1 Q-matrix.

Item	Initial rationale
1	A1 required to recognize that data are representations of people or things.
2	A1 required to count the five attributes shown as the table's columns.
3	A1 required to count the six observations shown as the table's rows.
4	A2 and A3 required to recognize n_floors is quantitative and select a density plot.
5	A4 required to identify percentages of category counts as relative frequencies.
6	A1 and A4 required to use the sample total 265 and compute 80/265 as about 30%.
7	A3 and A5 required to identify Q1 = 23 and use that 75% of values are at least Q1.
8	A5 required to compare mean and median and reject a left-skew claim when they are equal.
9	A5 required to place 10 well below Q1 = 23, so it is near the bottom few percent.
10	A3 required to read the middle 50% as the interval from Q1 to Q3.
11	A3 and A5 required to use mean 30 and SD 5 and apply the about 95% rule to get 20 to 40.
12	A4 and A6 required to recognize a conditional percent for categorical data is best shown with a bar chart.
13	A6 required to identify 64% as a column percent using the in-state total as the denominator.
14	A6 required to interpret Cramer's V = .89 as a strong relationship.
15	A8 required to identify the smoothed trend line as a line of averages.
16	A8 required to interpret Kendall's Tau of -0.62 and the plot as a moderate relationship.
17	A8 required to evaluate claims against the scatterplot and trend, leaving none of the statements supported.
18	A2 required to know scatterplots display two quantitative attributes.
19	A3 and A7 required to recognize grouped density curves as a side-by-side density plot.
20	A7 required to compare the two group distributions and identify the one with larger spread.
21	A7 required to judge the standardized difference between groups as about 1.0.
22	A7 required to know the ratio of standard deviations compares spread across groups.
23	A5 required to identify the distribution's shape as approximately normal.
24	A3 and A5 required to use the graph's center and near symmetry to estimate the mean at about 105 mg.
25	A3 and A5 required to read an approximately normal distribution and use the central 95% range of about 50 to 200.

Table 10: Initial version 1 rationales evaluated by the two experts.

Item	Refinement rationale
4	changed; A3 only since the task is to choose an appropriate graph for one quantitative variable.
6	changed; compute 80/265 as a percent from the frequency table, so A4 alone is needed.
7	changed; use Q1=23 from the five-number summary to judge the 75% statement, so A3 alone is needed.
12	changed; A4 only since the goal is to depict a category percent with a bar chart, and the suggested extra step is not needed.
13	unchanged; identify 64% as a column percent in a two-way table, so A6 only; RSS-suggested A5 is not used.
15	changed; A8 identifies the smoothed trend line, and A5 supports reading it as a line summarizing typical sleep values across screen time.
17	unchanged; check each claim directly against the scatterplot and trend line, so A8 only; the all-skill RSS alternative is not supported.
19	changed; the task is only to recognize the display as a density plot for a quantitative variable, so the graph-identification skill alone is needed.
20	unchanged; compare the spreads of the two quantitative group distributions, so A7 only; RSS-suggested A1 is not used.
25	changed; A5 only since the step is to identify the middle 95% range from the quantitative distribution.

Table 11: Version 1 refinement rationales evaluated by the two experts.

F Initial Q-Matrix Prompt

The following template was used for the initial Q-matrix proposal prompt. Application-specific item, attribute, and constraint blocks were inserted at the indicated locations. References (i.e., directories) to attached figures were added for items with visual content. Operational statistics items are not reproduced.

```
You are a subject-matter expert in [TEST DOMAIN] for [TARGET POPULATION].

TASK: Produce an initial Q-matrix for a cognitive
diagnostic model.

WORKFLOW (internal; do not print):
1) SOLVE: Determine a substantively defensible correct solution path for each item.
2) CODE: Set q_{jk}=1 if attribute k is required for successful solution of the
   item; else 0.
   - Do not weight by counts of operations.
   - Ignore non-attributes such as reading, formatting, and test-taking.
   - For polytomous items, code skills required for full credit.
   - Be parsimonious; avoid nice-to-have strategies.
3) SELF-CHECK: If the solution path and proposed Q disagree, fix Q internally.
4) Q-RATIONALE CONSISTENCY CHECK: Ensure each rationale names only attributes
   coded as 1. If not, revise the Q-vector before finalizing.

ITEMS (id: text with correct answer):
[ITEM BLOCK]

ATTRIBUTES (id: definition and examples):
[ATTRIBUTE BLOCK]

DOMAIN CONSTRAINTS (enforce):
[DOMAIN-CONSTRAINT BLOCK]
- Avoid all-zero or all-one columns unless strictly necessary.

[ATTACHED-FIGURE BLOCK, IF APPLICABLE]

OUTPUT (STRICT):
Return exactly TWO fenced R code blocks and nothing else.

FIRST R CODE BLOCK:
- Define `qmatrix` with exactly J rows and K columns using matrix(..., byrow=TRUE).
- Entries must be 0/1 only.
- Do not set row or column names.

SECOND R CODE BLOCK:
- Define `item_rationale` as a named list with exactly J entries named Item1, ...,
  ItemJ.
- Each entry must be one short character string.
- Give a step-based reason for the coded skills.
- Mention only skills coded as 1 and the key step that makes them necessary.
- Verify that each rationale matches its Q-vector.
- Do not explain every omitted skill.
- Keep each rationale short, concrete, and procedural.
```

G Q-Matrix Refinement Prompt

The following template was used after the initial fit. Only the subset of items flagged by RSS was inserted.

```
You are a subject-matter expert in [TEST DOMAIN] for [TARGET POPULATION].

TASK:
Revise the Q-matrix for THIS FLAGGED-ITEM SUBSET ONLY.

GOAL:
Use the statistical review only as diagnostic evidence. It may suggest possible
misspecification, but it is not the definitive standard. Final decisions must be
based primarily on item content and attribute definitions.

EVIDENCE PRIORITY:
```

```
1) Item content and a substantively defensible correct solution path
2) Attribute definitions
3) Statistical review of flagged items
4) Global structural/identifiability constraints

CURRENT Q FOR FLAGGED ITEMS:
[CURRENT-Q BLOCK]

STATISTICAL REVIEW OF FLAGGED ITEMS (advisory, not definitive):
[RSS-CRITIQUE BLOCK]

WORKFLOW (internal; do not print):
1) SOLVE: Determine a substantively defensible correct solution path for each
   flagged item.
2) CODE: Set q_{jk}=1 if attribute k is required for successful solution of the
   item; else 0.
   - Do not weight by counts of operations.
   - Ignore non-attributes such as reading, formatting, and test-taking.
   - For polytomous items, code skills required for full credit.
   - Be parsimonious; avoid nice-to-have strategies.
3) COMPARE the content-based Q-vector with the current vector and the statistically
   suggested competing vector.
4) DECIDE:
   - Treat the statistical review as diagnostic evidence only.
   - Adopt a suggested change only if supported by the item and attribute definitions.
   - If only part is supported, revise only that part.
   - If the current row is substantively correct, keep it.
   - Do not copy a suggested vector merely because it has better fit.
   - Change at least one flagged item. Choose the suggested alternative most
     consistent with content; if none is content-consistent, randomly apply one
     suggested alternative.
5) SELF-CHECK: If the solution path and proposed Q disagree, fix Q internally.
6) Q-RATIONALE CONSISTENCY CHECK: Ensure each rationale names only attributes coded
   as 1.

REVISION CONSTRAINTS:
- Revise only the flagged items shown here.
- Return a Q-vector for every flagged item in the same order as the current subset.
- Do not add or remove rows.
- Entries must be 0/1 only.
- Ensure at least one flagged item is changed.

ITEMS:
[FLAGGED-ITEM BLOCK]

ATTRIBUTES:
[ATTRIBUTE BLOCK]

DOMAIN CONSTRAINTS:
[DOMAIN-CONSTRAINT BLOCK]

[ATTACHED-FIGURE BLOCK, IF APPLICABLE]

OUTPUT (STRICT):
Return exactly TWO fenced R code blocks and nothing else.

FIRST R CODE BLOCK:
- Define `qmatrix` for the flagged subset using matrix(..., byrow=TRUE).
- Entries must be 0/1 only.
- Do not set row or column names.

SECOND R CODE BLOCK:
- Define `item_rationale` as a named list with one entry per flagged item.
- Begin each rationale with `changed;` or `unchanged;`.
- Give a brief, step-based reason tied to the solution path.
- If changed, state the step supporting the revised skills.
- If unchanged, state why the RSS suggestion was not supported by the solution path.
- Verify that every rationale matches the corresponding final Q-vector.
- Keep each rationale short, concrete, and procedural.
```

Composite Scores vs. Preference Rankings: Measuring Architecture Effects in LLM Feedback

Harvey Ngoe Kolle[1]
kolle@ualberta.ca

Carrie Demmans Epp[1,2]
cdemmansepp@ualberta.ca

Amna Liaqat[3]
aliaqat@gmu.edu

Maria Cutumisu[1,4,5]
maria.cutumisu@mcgill.ca

[1]Department of Computing Science, University of Alberta, Edmonton, AB, Canada
[2]Alberta Machine Intelligence Institute (Amii), Edmonton, AB, Canada
[3]George Mason University, Fairfax, VA, United States
[4]Department of Educational and Counselling Psychology, McGill University, Montreal, QC, Canada
[5]Mila – Quebec Artificial Intelligence Institute, Montreal, QC, Canada

Abstract

AI-generated feedback is often judged by quality ratings or preference rankings but rarely both. Comparing a multi-agent system, a single-agent system, and human feedback on student writing, we show that the two evaluation approaches can support different reported conclusions even when their underlying effects are nearly identical. These differences have consequences for how AI-generated feedback should be evaluated.

1 Background

Good feedback helps learners understand how their current performance compares with expected goals and what they need to do to improve (Hattie and Timperley, 2007; Shute, 2008). However, high quality feedback is hard to produce and is often sparse or delayed in practice (Gibbs and Simpson, 2005; Nicol and Macfarlane-Dick, 2006). Large language models (LLMs) offer one way to generate feedback at scale (Rüdian et al., 2025; Seßler et al., 2025; Lu et al., 2025) but the approaches can differ in quality. This creates a practical problem: how should competing methods for generating feedback be compared?

Most LLM systems generate feedback in a single step: one prompt has the model assess accuracy, pedagogical alignment, and tone all at once (Rüdian et al., 2025; Lu et al., 2025; Seßler et al., 2025), even though the feedback literature treats these as separate qualities (Hattie and Timperley, 2007; Shute, 2008). An alternative is to split generation into specialized agents, separating accuracy, pedagogical alignment, and tone (Kolle et al., 2026a; Cao et al., 2025). Whether this added complexity is worthwhile is an empirical question. It depends on whether the multi-agent approach produces feedback that raters judge as better than a single-agent system built from the same model and prompts, and on whether the comparison can detect such a difference at all.

However, whether such a difference can be detected depends on how feedback quality is measured, and current practice has two weaknesses. First, studies usually report either an aggregate quality score or a preference ranking (Rüdian et al., 2025; Weidlich et al., 2025) but rarely both, so it is unclear whether the two would point to the same answer. Second, the rating instruments often have unreported reliability (Nunnally, 1975; Streiner, 2003), which weakens confidence in any detected differences. Thus, choosing the better approach is partly a question of what counts as adequate evidence; the same concern that underlies argument-based approaches to validation in educational measurement (Kane, 2013). The present study addresses both problems using a 14-item, theory-grounded instrument that captures ten dimensions of feedback quality (Kolle et al., 2026b).

2 Aims

This study investigates whether the two evaluation approaches, composite quality scores and pairwise preference rankings, consistently indicate which feedback raters judge to be of higher quality. We examine this with three feedback conditions for adult English-language-learner writing: a multi-agent system, a single-agent system using the same base model, and human instructor feedback. The question matters because the choice of measure could determine whether one approach is better than another.

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 231–234
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

3 Sample

Eighteen raters (15 women, two men, and one non-binary) participated in this study. These raters were instructors and pre-service teachers with relevant experience evaluating writing or providing feedback. Age was available for 11 raters and averaged 25.9 years ($SD = 6.3$, range 19–40). Most raters reported prior experience with a teaching placement (72%), giving writing feedback (94%), grading writing (89%), feedback training (72%), and pedagogy or assessment training (78%). Fifteen raters (83%) had used LLMs before.

Student writing was from a prior study with adult English language learners from the Language Instruction for Newcomers to Canada program (Liaqat et al., 2019). Each rater evaluated three submissions and reviewed all three feedback versions for each submission, yielding 54 rated feedback instances. These 54 instances were drawn from 10 unique student submissions. Each submission was rated by between 2 and 8 raters across the three writing tasks.

4 Methods

We employed a within-subject design. The single-agent and multi-agent conditions used the same underlying language model (GPT-4.1) and task materials. Thus, the comparison isolated the effect of architecture rather than that of the model or task. The single-agent condition produced feedback in one step, using a single prompt that asked the model to address accuracy, pedagogical alignment, and tone. The multi-agent condition took that same initial output and passed it through a linear pipeline of specialized agents, each with its own prompt: successive agents checked grounding, corrected unsupported content, aligned feedback with task expectations, strengthened actionability, adjusted tone, and formatted the result. Both conditions shared the same base model and starting feedback, and differed only on whether that feedback was further processed by the pipeline. Full prompts and agent parameters are reported in the companion paper (Kolle et al., 2026a). The human condition used instructor-written comments from the original rubric-based feedback. For each submission, a rater reviewed three feedback versions. To reduce position effects, the version ordering was randomized and the source concealed.

The rater assessed each version using a 14-item instrument spanning ten dimensions of feedback quality (Table 1) and then ranked the three versions by overall preference. Composite scores, rescaled to 0–100, were analyzed with a linear mixed-effects model; condition was a fixed effect and rater and submission were random intercepts (Baayen et al., 2008). Preference rankings were analyzed with a Bradley-Terry model (Bradley and Terry, 1952). All pairwise comparisons were adjusted for multiple comparisons using Holm correction within each outcome.

5 Results

The 14-item instrument had high internal consistency, with $\omega = 0.917$, 95% CI $[0.903, 0.931]$, bootstrapped with $B = 2000$ resamples (Kolle et al., 2026b; McDonald, 2013), supporting its use as a composite outcome measure. Composite scores were highest for multi-agent feedback, followed by single-agent feedback and human feedback (Table 4). Both AI conditions scored higher than human feedback. The composite-score difference between multi-agent and single-agent feedback did not survive Holm correction (Table 2).

The ranking analysis showed the same ordering (Table 3): multi-agent feedback was preferred to single-agent feedback, and both AI conditions were preferred over human feedback. At the dimension level, the only significant difference between the multi-agent and single-agent conditions was on supportive tone, with multi-agent feedback rated higher ($\Delta = 4.63$, $p_{adj} = .049$, $d = 0.31$).

For the multi-agent vs. single-agent contrast, the two measures produced the same ordering and only marginally different effect sizes (composite $d = 0.29$, ranking $r = .31$); they differed only in whether the contrast reached significance after correction. A direct comparison confirmed that the measures themselves did not differ: the two feedback versions were ranked in the same direction as their composite scores on 72% of instances (95% CI $[0.62, 0.82]$), and the standardized effects of the two measures did not differ significantly (bootstrap difference $= 0.02$, 95% CI $[-0.32, 0.39]$).

6 Discussion and Conclusions

Composite scores and preference rankings agreed that both AI conditions were rated and ranked above human feedback. For the multi-agent vs. single-agent contrast, they agreed on ordering and effect magnitude but differed on whether the contrast reached significance. This difference in signif-

Table 1: The 14-item instrument used to rate feedback quality, grouped into ten dimensions. Items marked [R] are reverse-scored before summing so that the composite consistently reflects higher perceived quality.

Dimension	**Item**
Specific & Clear	This feedback targets specific parts of the writing. This feedback is clearly written.
Goal-Referenced	This feedback explains how the writing meets the expected goals.
Actionable	This feedback offers clear guidance the student can follow. This feedback supports improvement in writing skills.
Supportive Tone	This feedback maintains a positive tone. This feedback is discouraging.[R]
Trustworthiness	This feedback is inaccurate.[R]
Objective	This feedback is based on evidence from the student's writing. This feedback focuses on student traits.[R]
Non-Threatening	This feedback uses harsh language.[R]
Non-Revealing	This feedback suggests alternative phrasing.
Conciseness	This feedback avoids overly detailed elaboration.
Overall	This feedback would be appropriate to share directly with a student.

Table 2: Pairwise contrasts for composite feedback quality scores. 95% CIs are unadjusted; p_{adj} values are Holm-corrected.

Comparison	Diff	95% CI	p_{adj}	d
Multi-Agent vs. Single-Agent	2.46	[0.14, 4.78]	.114	0.29
Multi-Agent vs. Human	26.38	[22.48, 30.27]	$< .001$	1.85
Single-Agent vs. Human	23.92	[20.24, 27.59]	$< .001$	1.78

Table 3: Pairwise preference probabilities from the Bradley-Terry model.

Comparison	P(preferred)	p_{adj}	r
Multi-Agent vs. Single-Agent	.657	.015	.31
Multi-Agent vs. Human	.957	$< .001$	.91
Single-Agent vs. Human	.921	$< .001$	.84

Table 4: Descriptive statistics for composite scores by condition (rescaled to 0–100).

Condition	M	SD	Range
Multi-Agent	87.0	11.2	54–100
Single-Agent	84.5	12.6	51–100
Human	60.6	13.5	40–91

icance is not itself evidence that the measures differ (Gelman and Stern, 2006), and a direct comparison found no significant difference between them.

The pattern of both AI conditions outperforming human feedback is consistent with prior reports that LLM-generated feedback can compare favorably with human feedback on some dimensions (Rüdian et al., 2025; Weidlich et al., 2025). However, the human feedback in this study was drawn from a prior instructional context, so the observed difference should be interpreted cautiously.

The split in significance between the two AI conditions may reflect that a difference concentrated in one dimension (i.e., tone) is diluted once combined into a 14-item composite (Nunnally, 1975). However, with 18 raters, it may reflect that the ranking analysis was better powered to detect small effects. A larger-sample replication is planned to tease apart these explanations. If the tone difference persists, it would be consistent with the multi-agent pipeline's design that uses a dedicated tone agent to revise feedback after its content and pedagogical structure have already been refined (Kolle et al., 2026a).

These results suggest that the choice between aggregate scores and preference rankings is not just a reporting detail: it can determine whether a difference between AI systems is detected at all. More broadly, this implies that conclusions

about AI-generated feedback are partly an artifact of how feedback quality is measured, not just of the systems themselves. As these tools become more common, evaluations that only report one outcome measure run the risk of over- or under-stating real differences between systems. Therefore, comparisons of AI systems should report both where possible.

References

R.H. Baayen, D.J. Davidson, and D.M. Bates. 2008. Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, 59(4):390–412. Special Issue: Emerging Data Analysis.

Ralph Allan Bradley and Milton E Terry. 1952. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345.

Jie Cao, Chloe Qianhui Zhao, Xian Chen, Shuman Wang, Christian Schunn, Kenneth R Koedinger, and Jionghao Lin. 2025. From first draft to final insight: A multi-agent approach for feedback generation. In *International Conference on Artificial Intelligence in Education*, pages 163–176. Springer-Verlag.

Andrew Gelman and Hal Stern. 2006. The difference between "significant" and "not significant" is not itself statistically significant. *The American Statistician*, 60(4):328–331.

Graham Gibbs and Claire Simpson. 2005. Conditions under which assessment supports students' learning. *Learning and Teaching in Higher Education*, (1):3–31.

John Hattie and Helen Timperley. 2007. The power of feedback. *Review of educational research*, 77(1):81–112.

Michael T. Kane. 2013. Validating the interpretations and uses of test scores. *Journal of Educational Measurement*, 50(1):1–73.

Harvey Ngoe Kolle, Carrie Demmans Epp, Amna Liaqat, and Maria Cutumisu. 2026a. Do Language Models Matter? Evaluating Model Choice in a Multi-agent Feedback Framework. In *Artificial Intelligence in Education. Late Breaking Results, WideAIED, Practitioners, Industry and Policies, Blue Sky, Doctoral Consortium, FoL Workshops and Tutorials, FoL Invited Papers*, volume 3031, pages 108–114, Cham. Springer Nature Switzerland. Series Title: Communications in Computer and Information Science.

Harvey Ngoe Kolle, Carrie Demmans Epp, Amna Liaqat, and Maria Cutumisu. 2026b. Evaluating the pedagogical quality of LLM-generated feedback: A criterion-based and comparative study. In *Proceedings of the First International Workshop on Pedagogical Evaluation of Automated Feedback (PEAF 2026)*, CEUR Workshop Proceedings. CEUR-WS.org.

Amna Liaqat, Gokce Akcayir, Carrie Demmans Epp, and Cosmin Munteanu. 2019. Mature ELLs' Perceptions Towards Automated and Peer Writing Feedback. In *Transforming Learning with Meaningful Technologies*, Lecture Notes in Computer Science, pages 266–279. Springer International Publishing.

Xinyi Lu, Aditya Mahesh, Zejia Shen, Mitchell Dudley, Larissa Sano, and Xu Wang. 2025. Exploring LLM-generated feedback for economics essays: How teaching assistants evaluate and envision its use. In *International Conference on Artificial Intelligence in Education*, pages 392–406, Berlin, Heidelberg. Springer-Verlag.

Roderick P McDonald. 2013. *Test theory: A Unified Treatment*. Psychology Press.

David J. Nicol and Debra Macfarlane-Dick. 2006. Formative assessment and self-regulated learning: A model and seven principles of good feedback practice. *Studies in Higher Education*, 31(2):199–218.

Jum C Nunnally. 1975. Psychometric theory—25 years ago and now. *Educational researcher*, 4(10):7–21.

Sylvio Rüdian, Julia Podelo, Jakub Kužílek, and Niels Pinkwart. 2025. Feedback on feedback: Students' perceptions for feedback from teachers and few-shot LLMs. In *Proceedings of the 15th International Learning Analytics and Knowledge Conference*, pages 82–92, New York, NY, USA. Association for Computing Machinery.

Kathrin Seßler, Maurice Fürstenberg, Babette Bühler, and Enkelejda Kasneci. 2025. Can AI grade your essays? a comparative analysis of large language models and teacher ratings in multidimensional essay scoring. In *Proceedings of the 15th International Learning Analytics and Knowledge Conference*, pages 462–472, New York, NY, USA. Association for Computing Machinery.

Valerie J. Shute. 2008. Focus on formative feedback. *Review of Educational Research*, 78(1):153–189.

David L. Streiner. 2003. Starting at the beginning: An introduction to coefficient alpha and internal consistency. *Journal of Personality Assessment*, 80(1):99–103.

Joshua Weidlich, Flurin Gotsch, Kai Schudel, Claudia Marusic-Würscher, Jennifer Mazzarella, Hannah Bolten, Dario Bütler, Simon Luger, Bettina Wohlfehnder, and Katharina Maag Merki. 2025. Teacher, peer, or AI? comparing effects of feedback sources in higher education. *Computers and Education Open*, page 100300.

Identifying Communication Profiles from Physician–Patient Conversations Using Large Language Models

Reyhaneh Hosseinpourkhoshkbari * **Yash Bipin Jain** **Richard M. Golden**

School of Behavioral and Brain Sciences, The University of Texas at Dallas

Abstract

Large language models (LLMs) are increasingly used to score clinical communication, but validation is often limited to individual checklist items or total scores. These measures may not capture how communication behaviors occur together within a transcript. We therefore examine *communication profiles*: recurring combinations of behaviors that characterize different patterns of clinician communication and may support more targeted formative feedback. We analyzed 213 simulated respiratory OSCE transcripts rated on 15 binary Kalamazoo-derived communication items and compared final human ratings with GPT-4o, GPT-o3, and GPT-5.6 Sol. Raw item-level agreement was relatively high overall, but varied substantially across behaviors and was lower for several judgment-intensive items used for profile modeling. Bayesian latent-class analysis identified three stable human-derived profiles. Although the LLM-derived profiles showed broadly similar item-probability patterns, the models frequently assigned individual transcripts to different profiles than the human ratings. These findings show that agreement on individual communication skills does not necessarily translate into agreement on higher-level communication profiles. If LLMs are used to provide profile-based feedback, validation should therefore include profile-level agreement in addition to item-level performance.

1 Introduction

Clinical communication is central to safe and effective care. Communication quality is associated with patient satisfaction, adherence, recovery, and other clinical outcomes, while communication failures can contribute to errors and poorer care (Chavda et al., 2020; Hosseinpourkhoshkbari and Golden, 2025; Singh et al., 2013; Ahmadipour and Mortazavi, 2015). Simulated encounters and standardized patients provide controlled settings in which these skills can be practiced and assessed (Hosseinpourkhoshkbari and Golden, 2025; Mavis et al., 2006; Johnson et al., 2020). The Kalamazoo Essential Elements framework provides a widely used structure for evaluating communication behaviors across a transcript (Makoul, 2001; Dalus et al., 2024; Yoon and Michaelsen, 2015). Because assessment outcomes can vary across evaluators, establishing and reporting the reliability of these ratings is essential.

Communication assessments are often summarized using total or domain-level scores, although reliability and agreement can vary substantially across individual communication domains (Fan et al., 2024). Examining item-level performance can identify specific strengths and weaknesses that may be hidden by similar overall scores and can therefore provide more informative diagnostic feedback (Ark et al., 2024). Moreover, two transcripts with similar total scores may show different combinations of communication strengths and weaknesses. These joint patterns may be particularly important for formative feedback.

Communication assessment is also difficult to perform consistently at scale (Brannick et al., 2011). Human raters tend to agree more readily on directly observable behaviors than on behaviors requiring more subjective judgments, such as empathy or understanding the patient's perspective (Roh et al., 2017). LLMs have recently shown promise for automating assessment of medical interviews and OSCE performance (Shakur et al., 2024; Geathers et al., 2025), but most validation focuses on individual items or aggregate scores.

As a result, we examine communication profiles: recurring patterns in how multiple communication behaviors co-occur within a transcript. If an automated system is ultimately used to characterize these broader patterns for formative feedback, the relevant question is not only whether individ-

*reyhaneh.hosseinpour@utdallas.edu

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 235–247
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

ual items are scored similarly to human raters, but whether the same higher-level profiles are recovered.

We therefore distinguish two levels of agreement. *Item-level agreement* asks whether human and LLM raters make the same judgment on individual communication behaviors. *Profile-level agreement* asks whether they assign the same learner to the same multidimensional communication pattern. Our central question is:

> *Does Item-level agreement on communication behaviors translate into Profile-level agreement?*

We make four contributions. First, we compare human ratings with GPT-4o, GPT-o3, and GPT-5.6 Sol across all 15 Kalamazoo-derived items. Second, we compare the three LLMs item by item on the eight informative behaviors used for profile modeling, highlighting model-specific disagreement patterns. Third, we identify latent communication profiles using Bayesian latent-class modeling and evaluate their stability with resampling and an alternative clustering method. Fourth, we compare human-derived profiles with profiles obtained from all three LLMs.

2 Data and Annotation

2.1 Dataset

We used the respiratory-case subset of the simulated patient–physician interview dataset of Fareez et al. (2022). The corpus contains 213 OSCE-style interviews, each approximately 50 conversational turns, conducted by senior medical students and resident physicians in respiratory history-taking scenarios.

The human annotations were produced by three coders using the same operational definitions. A 50-transcript holdout set was independently rated by all three coders to assess inter-rater agreement. Disagreements were discussed and resolved to obtain the final item ratings. Complete final human ratings and ratings from GPT-4o, GPT-o3, and GPT-5.6 Sol were available for all 213 transcripts, which formed the paired comparison sample.

2.2 Assessment Instrument

Communication was coded using 15 binary Kalamazoo-derived behaviors (Table 1). Each rater reviewed the complete transcript and marked whether each behavior was present. The same operational definitions were used for human and LLM ratings; the full scoring prompt is included in Appendix A.

Item	Operationalized behavior
Q1	Greet the patient
Q2	Use words showing care and concern
Q3	Allow the opening statement without interruption
Q4	Ask “Is there anything else?” or equivalent
Q5	Explain or negotiate an agenda
Q6	Begin with the patient’s story using open questions
Q7	Clarify details with specific questions
Q8	Summarize and invite correction/addition
Q9	Transition effectively to additional questions
Q10	Ask about life circumstances affecting health
Q11	Elicit beliefs, concerns, and expectations
Q12	Respond explicitly to feelings or values
Q13	Explain using words easy for the patient to understand
Q14	Ask whether the patient has questions or other issues
Q15	Acknowledge the patient and close the interview

Table 1: The 15 Kalamazoo-derived communication behaviors used in the analysis.

3 Methods

3.1 Agreement Analysis

Several binary items had highly imbalanced endorsement rates. Under this condition, Cohen’s κ can be strongly affected by prevalence because its chance-agreement correction can become large when raters frequently use the same category. We therefore used raw percent agreement and Gwet’s AC1 (Gwet, 2008). Gwet’s AC1 is a chance-corrected agreement coefficient that is less sensitive to extreme category prevalence, making it more appropriate for this rubric-based setting, where highly imbalanced response categories can produce counterintuitive chance-corrected agreement estimates despite high observed agreement. Gwet’s AC1 was interpreted as a continuous measure of chance-corrected agreement, with larger values indicating stronger agreement.

For the 50 independently rated holdout transcripts, human inter-rater agreement was summarized using mean pairwise raw agreement and multi-rater Gwet’s AC1. LLM–human agreement was calculated against the final human ratings on the full paired sample of 213 transcripts.

3.2 Selection of Informative Items

Profiles were estimated from items that showed sufficient variation across the final human ratings. Items that were endorsed by humans in fewer than 20% or more than 90% of transcripts provided little information for distinguishing one transcript from another, so they were excluded from the profile analysis. These items were Q1, Q3, Q5, Q6, Q7, Q9, and Q10. They were still included in the item-level agreement analysis. The remaining eight items—Q2, Q4, Q8, Q11, Q12, Q13, Q14, and Q15—were used to estimate the communication profiles.

3.3 Bayesian Latent Class Analysis

We used Bayesian latent-class analysis to identify recurring communication patterns across transcripts rather than reducing the ratings to a single total score (Li et al., 2018). The Bayesian approach allows uncertainty in the profile prevalences and profile-specific item probabilities to be represented through posterior distributions rather than treating these quantities as fixed estimates. Two transcripts may have the same number of endorsed items but differ in which items are present. The analysis included $N = 213$ transcripts and $J = 8$ binary communication items. Let $Y_{ij} = 1$ if item j is scored as present in transcript i, and $Y_{ij} = 0$ otherwise.

Latent profile structure. We assume that each transcript belongs to one of K latent communication profiles. These profiles are not observed directly but are inferred from recurring patterns across the binary communication items. Let $Z_i \in \{1, \ldots, K\}$ denote the latent profile of transcript i. The profile proportions are

$$\boldsymbol{\pi} = (\pi_1, \ldots, \pi_K), \qquad \sum_{k=1}^{K} \pi_k = 1.$$

The latent profile indicator is modeled as

$$Z_i \mid \boldsymbol{\pi} \sim \text{Categorical}(\boldsymbol{\pi}), \qquad \Pr(Z_i = k) = \pi_k,$$

where π_k represents the prevalence of profile k, that is, how common that profile is across transcripts.

Profile-specific item probabilities. Each profile is defined by the probability that each communication item is scored as present. Let

$$\theta_{kj} = \Pr(Y_{ij} = 1 \mid Z_i = k),$$

where θ_{kj} is the probability that item j is scored as present for a transcript belonging to profile k. Because the items are binary,

$$\Pr(Y_{ij} \mid Z_i = k) = \theta_{kj}^{Y_{ij}} (1 - \theta_{kj})^{1 - Y_{ij}}.$$

The complete set of item probabilities defining profile k is represented by the profile-specific item-probability vector

$$\boldsymbol{\theta}_k = (\theta_{k1}, \ldots, \theta_{kJ}).$$

Thus, each element of $\boldsymbol{\theta}_k$ represents the estimated probability that a particular communication item is scored as present for transcripts belonging to that profile. Throughout the paper, comparisons of human- and LLM-derived profiles are based on these profile-specific item-probability vectors.

Latent-class analysis assumes conditional independence of the items given the profile. Thus, for $\mathbf{Y}_i = (Y_{i1}, \ldots, Y_{iJ})$,

$$\Pr(\mathbf{Y}_i \mid Z_i = k) = \prod_{j=1}^{J} \theta_{kj}^{Y_{ij}} (1 - \theta_{kj})^{1 - Y_{ij}}.$$

This expression gives the probability of transcript i's complete item pattern under profile k.

Bayesian priors. We specified two priors. First, the profile proportions were assigned a symmetric Dirichlet prior,

$$\boldsymbol{\pi} \sim \text{Dirichlet}(2, \ldots, 2),$$

where

$$p(\boldsymbol{\pi}) = \frac{\Gamma(2K)}{\Gamma(2)^K} \prod_{k=1}^{K} \pi_k, \qquad \sum_{k=1}^{K} \pi_k = 1.$$

This prior mildly discourages near-empty profiles while treating all profile labels symmetrically.

Second, each profile-specific item probability was assigned

$$\theta_{kj} \sim \text{Beta}(1, 1),$$

where

$$p(\theta_{kj}) = \frac{\theta_{kj}^0 (1 - \theta_{kj})^0}{B(1, 1)} = 1, \qquad 0 \le \theta_{kj} \le 1.$$

Thus, the item-probability prior is uniform on $[0, 1]$ and does not favor high or low endorsement probabilities.

Bayesian posterior. Let $\boldsymbol{\Theta} = \{\theta_{kj}\}$ denote the complete set of profile-specific item probabilities. The posterior distribution of $\boldsymbol{\pi}$ and $\boldsymbol{\Theta}$ given the observed data $\mathbf{Y}$ is

$$p(\boldsymbol{\pi}, \boldsymbol{\Theta} \mid \mathbf{Y}) \propto p(\mathbf{Y} \mid \boldsymbol{\pi}, \boldsymbol{\Theta})\, p(\boldsymbol{\pi})\, p(\boldsymbol{\Theta}).$$

The likelihood is obtained by marginalizing over the unknown latent profile indicator for each transcript:

$$p(\mathbf{Y} \mid \boldsymbol{\pi}, \boldsymbol{\Theta}) = \prod_{i=1}^{N} \left[\sum_{k=1}^{K} \pi_k \prod_{j=1}^{J} \theta_{kj}^{Y_{ij}} (1-\theta_{kj})^{1-Y_{ij}} \right],$$

with

$$p(\boldsymbol{\Theta}) = \prod_{k=1}^{K} \prod_{j=1}^{J} p(\theta_{kj}).$$

Posterior profile probabilities and classification. For a given posterior draw of the model parameters, the conditional probability that transcript i belongs to profile k is

$$\tau_{ik} = \Pr(Z_i = k \mid \mathbf{Y}_i, \boldsymbol{\pi}, \boldsymbol{\Theta}) = \frac{s_{ik}}{\sum_{\ell=1}^{K} s_{i\ell}},$$

where

$$s_{ik} = \pi_k \prod_{j=1}^{J} \theta_{kj}^{Y_{ij}} (1-\theta_{kj})^{1-Y_{ij}}.$$

These probabilities were computed at each posterior draw and then averaged across draws. Let

$$\bar{\tau}_{ik} = \frac{1}{S} \sum_{s=1}^{S} \tau_{ik}^{(s)}$$

denote the posterior-averaged probability that transcript i belongs to profile k, where S is the total number of retained posterior draws. We refer to $\bar{\tau}_{ik}$ as the *posterior profile probability*. These posterior profile probabilities reflect both how well a transcript's item pattern fits each profile and uncertainty in the estimated model parameters. Each transcript was placed into the profile with the largest posterior profile probability,

$$\hat{Z}_i = \arg\max_k \bar{\tau}_{ik}.$$

For each transcript, the *maximum posterior profile probability*,

$$\max_k \bar{\tau}_{ik},$$

was used as a measure of classification certainty. Values closer to 1 indicate that one profile is clearly favored, whereas lower values indicate greater uncertainty across profiles.

Model estimation and profile selection. Models were estimated using Markov Chain Monte Carlo (MCMC) with the No-U-Turn Sampler (NUTS) (Hoffman and Gelman, 2014), implemented in PyMC (Salvatier et al., 2016). MCMC was used to generate samples from the posterior distribution of the model parameters, and NUTS provided an efficient sampling algorithm for these continuous parameters. We used four chains, 2,000 tuning iterations, 1,000 retained posterior draws per chain, and a target acceptance rate of 0.95. We fit $K = 2, 3,$ and 4 solutions and compared AIC, BIC, and MCMC convergence using $\hat{R}$, while also examining whether the resulting profiles were non-negligible and interpretable. $\hat{R}$ assesses convergence across MCMC chains, with values close to 1 indicating that the chains sampled from the same posterior distribution.

Applying model to LLM ratings. After selecting K from the final human ratings, we fit the same eight-item model separately to GPT-4o, GPT-o3, and GPT-5.6 Sol. Because latent-class labels are arbitrary, each LLM-derived profile was matched to the human-derived profiles by maximizing cosine similarity between the corresponding profile-specific item-probability vectors $\boldsymbol{\theta}_k$. This matching step only relabels the profiles for comparison; it does not alter the item ratings or the transcript-level profile probabilities.

3.4 Robustness and Complementary Analyses

We assessed stability using 200 bootstrap resamples and 200 refits after randomly dropping 5% of transcripts. We also applied PAM (k-medoids) clustering with Jaccard distance to the same eight binary items as a non-mixture comparison. In contrast to latent-class analysis, PAM groups transcripts based on similarity without assuming an underlying probabilistic mixture model.

4 Results

4.1 Agreement Varies Across Items and Evaluator Types

Among the three human coders, mean pairwise raw agreement across all 15 items was 82.6%, and mean multi-rater Gwet's AC1 was 0.712. For the

eight profile-modeling items retained based on variation in the final human ratings, mean pairwise raw agreement decreased to 70.5%, and mean multirater Gwet's AC1 decreased to 0.498. Thus, human agreement was lower for the eight items used to construct the communication profiles.

LLM–human agreement showed a similar pattern. Across all 15 items, mean Gwet's AC1 was 0.773 for GPT-4o, 0.718 for GPT-o3, and 0.701 for GPT-5.6 Sol. Across the eight profile-modeling items, mean Gwet's AC1 decreased to 0.600, 0.514, and 0.480, respectively. On the full 213-transcript paired sample, mean raw agreement across the eight items was 74.4% for GPT-4o, 71.2% for GPT-o3, and 71.0% for GPT-5.6 Sol (Table 2).

LLM–human agreement varied substantially across communication behaviors (Figure 1). Agreement was consistently high for Q1, Q3, Q5–Q7, Q9–Q10, Q13, and Q14, while Q2 showed consistently lower agreement across all three models. Q8, Q11, and Q12 showed the largest model-dependent differences, with particularly low agreement for some models.

Overall, agreement patterns were qualitatively similar across human–human, human–LLM, and LLM–LLM comparisons, with consistently lower agreement for Q2, Q8, Q11, and Q12, although the magnitude of disagreement varied across evaluator types. Across the eight profile-modeling items, mean Gwet's AC1 was 0.498 among the human coders and 0.600, 0.514, and 0.480 for GPT-4o, GPT-o3, and GPT-5.6 Sol versus the final human ratings, respectively; under the Landis and Koch (1977) benchmarks, these values fall within the moderate agreement range.

4.2 Bayesian Latent Class Model Selection Results

Table 3 summarizes model selection using the Bayesian Latent Class Analysis. The $K = 3$ solution had the lowest AIC and BIC. All models converged without divergent transitions; for $K = 3$, the maximum $\hat{R}$ for the profile proportions was 1.002. We therefore retained $K = 3$ as the primary solution.

The three profiles showed distinct patterns (Table 4; Figure 2). Profile 1 was characterized by high probabilities for Q4, Q13, and Q14, moderate probabilities for Q15 and Q8, and low probabilities for Q2, Q11, and Q12. Profile 2 showed the broadest pattern of high probabilities, including Q4, Q11, Q12, Q13, Q14, and Q15, with moderate Q8 and low Q2. In contrast, Profile 3 showed high probabilities for Q2, Q13, Q14, and Q15, but low probabilities for Q4, Q8, Q11, and Q12. The estimated profile prevalences were 44%, 34%, and 23%, respectively.

The distribution of maximum posterior profile probabilities showed variation in classification certainty (Figure 3). Many transcripts were assigned to a profile with very high posterior probability, whereas a smaller group showed more moderate probabilities around 0.50–0.65, indicating greater uncertainty between profiles. Very few transcripts had probabilities close to the theoretical minimum of $1/3$, suggesting that highly ambiguous profile classification were uncommon (Figure 3, A).In addition, most transcripts were clearly associated with one profile, while a smaller subset showed overlap across profiles, indicating greater classification uncertainty (Figure 3, B).

4.3 The Human-Derived Profiles Are Stable

The $K = 3$ solution was recovered in 99% of 200 bootstrap resamples and 100% of 200 refits after randomly dropping 5% of transcripts. PAM clustering with Jaccard distance produced a complementary partition with moderate similarity to the latent-class solution (Adjusted Rand Index [ARI] = 0.56).

4.4 LLM-Derived Profiles Show Overall Pattern Similarity but Item-Specific Differences

After matching the LLM-derived profiles to the human-derived profiles, cosine similarity between the corresponding profile-specific item-probability vectors ranged from .81 to .99 (Figure 4). These values indicate similarity in the overall direction of the eight-item profile patterns, but they do not imply close agreement for every individual item.

The degree of item-probability correspondence varied across profiles. For Profile 1, the human and LLM solutions were similar for several items, including Q4, Q11, and Q12, but differed substantially for Q13 and Q14. For Profile 3, similar probabilities were observed for Q13–Q15, whereas large differences were evident for Q2 and Q4. Profile 2 showed the strongest overall correspondence, although differences remained for Q2, Q8, and Q11. Thus, the LLM-derived profiles captured some broad features of the human-derived patterns while differing substantially on particular communication behaviors.

Item	GPT-4o		GPT-o3		GPT-5.6 Sol	
	Raw	AC1	Raw	AC1	Raw	AC1
Q2 Care	62.4%	.352	67.1%	.371	68.5%	.372
Q4 Full concerns	78.4%	.733	78.4%	.733	78.4%	.733
Q8 Summarize	73.2%	.552	69.5%	.417	54.5%	.093
Q11 Elicit	62.9%	.397	61.0%	.352	61.0%	.221
Q12 Respond	66.7%	.431	54.5%	.090	53.5%	.071
Q13 Explain	85.9%	.836	78.4%	.725	84.5%	.807
Q14 Ask questions	89.7%	.838	84.5%	.744	91.5%	.871
Q15 Close	76.1%	.660	76.5%	.679	75.6%	.675
Mean	**74.4%**	**.600**	**71.2%**	**.514**	**71.0%**	**.480**

Table 2: LLM–human agreement for the eight profile-modeling items on the full paired sample ($N = 213$). Raw agreement and Gwet's AC1 are calculated against the final human ratings. The independent 50-transcript holdout is used separately for human inter-rater agreement.

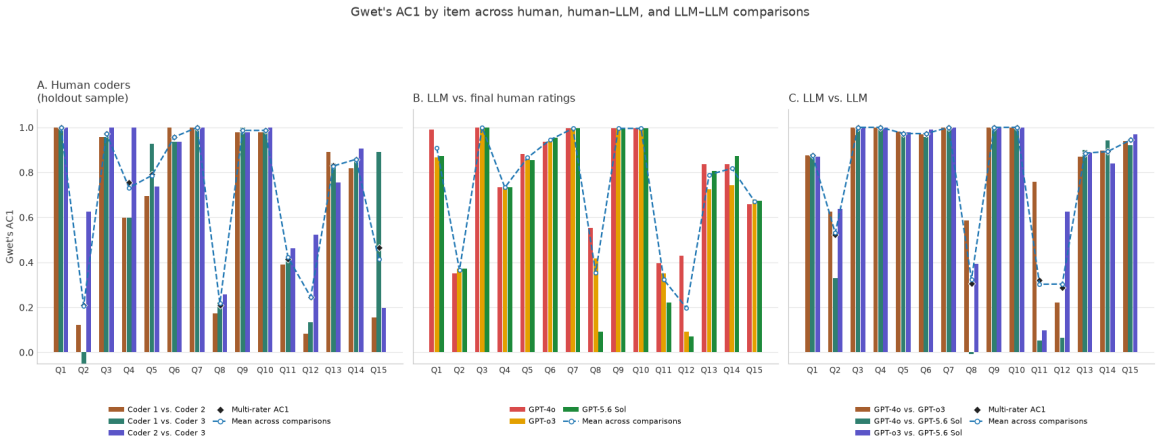


Figure 1: Item-level Gwet's AC1 across evaluator types. Panel A shows pairwise and multi-rater agreement among the three human coders on the 50 independently rated holdout transcripts. Panel B shows each LLM against the final human ratings on all 213 transcripts. Panel C shows pairwise and multi-rater agreement among the three LLMs on all 213 transcripts. Overall, agreement patterns were qualitatively similar across human–human, human–LLM, and LLM–LLM comparisons, with consistently lower agreement for Q2, Q8, Q11, and Q12, although the magnitude of disagreement varied across evaluator types.

Profiles	AIC	BIC
$K = 2$	1897.9	1955.0
$K = 3$	**1861.6**	**1949.0**
$K = 4$	1863.2	1980.8

Table 3: Model-selection statistics for the human-derived eight-item latent-class models. Lower values indicate better fit after penalizing model complexity.

Profile prevalence also differed across the human and LLM solutions (Table 5). The GPT-4o solution estimated a substantially higher prevalence for Profile 3, whereas GPT-o3 and GPT-5.6 Sol estimated higher prevalences for Profile 2. Profile 1 had a lower estimated prevalence in all three LLM solutions (6–10%) than in the human solution (44%). The prevalences reported in Table 5 represent the model-estimated profile proportions, π_k, and therefore need not equal the proportions obtained by classifying each transcript according to its largest posterior profile probability.

Transcript-level profile classification agreement with the human ratings was 29.1% for GPT-4o, 39.0% for GPT-o3, and 38.5% for GPT-5.6 Sol. In contrast, GPT-o3 and GPT-5.6 Sol agreed with each other on 77.9% of transcript classifications (Table 6).

Figure 5 shows substantial differences in transcript-level profile classifications between the human- and LLM-derived solutions. Among the human-derived classifications, 93 transcripts were classified as Profile 1. In contrast, only 3 transcripts were classified as Profile 1 by GPT-4o, 15 by GPT-o3, and 10 by GPT-5.6 Sol. GPT-4o instead classified most transcripts as Profile 3 ($n = 156$), whereas GPT-o3 and GPT-5.6 Sol classified most

Profile	High ($\geq$ .70)	Moderate (.30–.69)	Low ($<$ .30)	Prev.
Profile 1	Q4 (.93), Q13 (.89), Q14 (.70)	Q15 (.59), Q8 (.37)	Q2 (.25), Q12 (.14), Q11 (.11)	44%
Profile 2	Q4 (.96), Q14 (.86), Q12 (.82), Q13 (.79), Q15 (.74), Q11 (.73)	Q8 (.40)	Q2 (.29)	34%
Profile 3	Q15 (.97), Q13 (.93), Q2 (.93), Q14 (.76)	—	Q11 (.29), Q8 (.26), Q4 (.22), Q12 (.04)	23%

Table 4: Profile-specific item-probability patterns for the three human-derived communication profiles. High, moderate, and low refer to estimated item probabilities of $\geq$ .70, .30–.69, and $<$ .30, respectively. Prevalence represents the model-estimated proportion of transcripts in each profile.

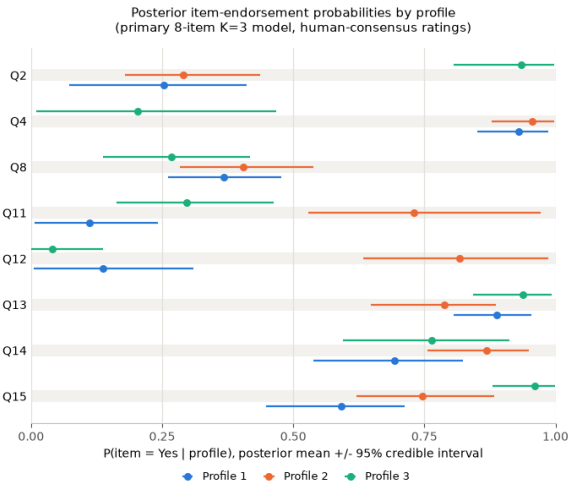


Figure 2: Posterior item-endorsement probabilities for the human-derived $K = 3$ model. Points show posterior means and intervals show 95% credible intervals.

Source	Profile 1	Profile 2	Profile 3
Human	44%	34%	23%
GPT-4o	9%	24%	67%
GPT-o3	10%	55%	35%
GPT-5.6 Sol	6%	65%	29%

Table 5: Model-estimated prevalence of the matched profiles ($N = 213$).

transcripts as Profile 2 ($n = 134$ and $n = 139$, respectively).

5 Discussion

This study examined LLM-based clinical communication assessment at both the item and profile levels. Overall item-level agreement between the LLMs and the final human ratings was generally comparable to agreement among human coders, but disagreement was concentrated in particular behaviors, especially Q2, Q8, Q11, and Q12. Bayesian latent-class analysis identified three distinct and stable communication profiles from the human ratings. The LLM-derived profiles showed overall

Pair	Profile agreement
Human–GPT-4o	29.1%
Human–GPT-o3	39.0%
Human–GPT-5.6 Sol	38.5%
GPT-4o–GPT-o3	54.0%
GPT-4o–GPT-5.6 Sol	51.6%
GPT-o3–GPT-5.6 Sol	**77.9%**

Table 6: Raw agreement on matched transcript-level profile classifications ($N = 213$).

pattern similarity to the human-derived profiles, but substantial differences remained for specific items. Profile prevalence also differed across human and LLM solutions, and transcript-level profile classification agreement was relatively low. Thus, similarity in the overall profile pattern did not necessarily translate into similarity in individual item probabilities, profile prevalence, or transcript-level classification. Neither GPT-o3 nor GPT-5.6 Sol showed a consistent improvement over GPT-4o across these measures.

These findings have several implications for evaluating LLM-based communication assessment. Overall agreement alone can be misleading because it may hide substantial disagreement on particular communication behaviors. Item-level analyses are therefore needed to identify where human and LLM ratings agree and where they differ.

Evaluation should also match the level at which the assessment will be interpreted. If an LLM is used to score individual checklist items, item-level agreement may be the primary concern. However, if those ratings are used to identify communication profiles or provide profile-based feedback, the resulting profile patterns, prevalences, and transcript-level classifications should also be validated.

Profile-based methods such as Bayesian latent-class analysis can identify recurring combinations of communication behaviors that may not be visible from total scores alone. However, before such

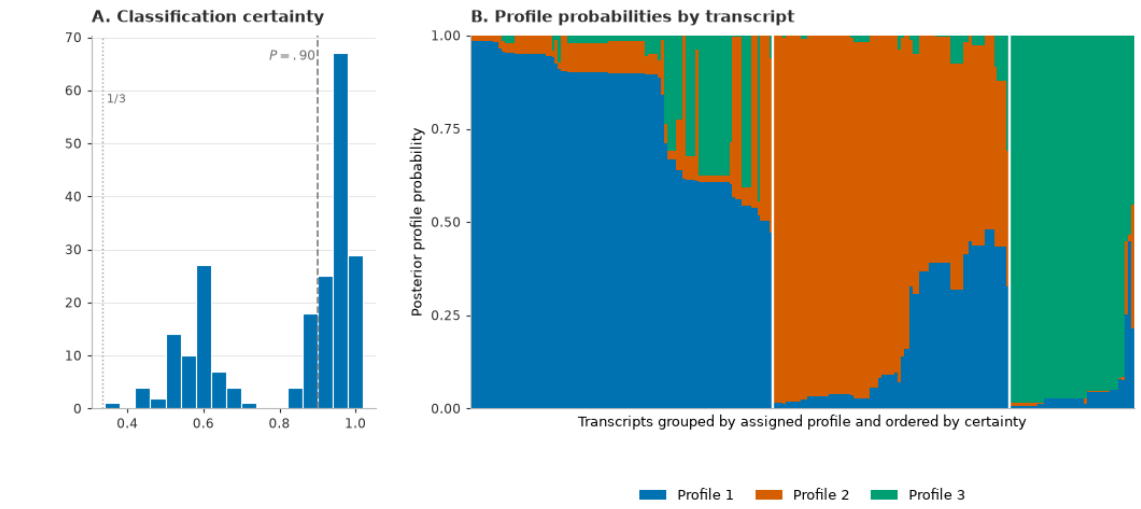


Figure 3: Posterior classification certainty for the human-derived $K = 3$ model ($N = 213$). Panel A shows each transcript's maximum posterior profile probability. Panel B shows posterior profile probabilities, with transcripts grouped by assigned profile and sorted within each profile by certainty.

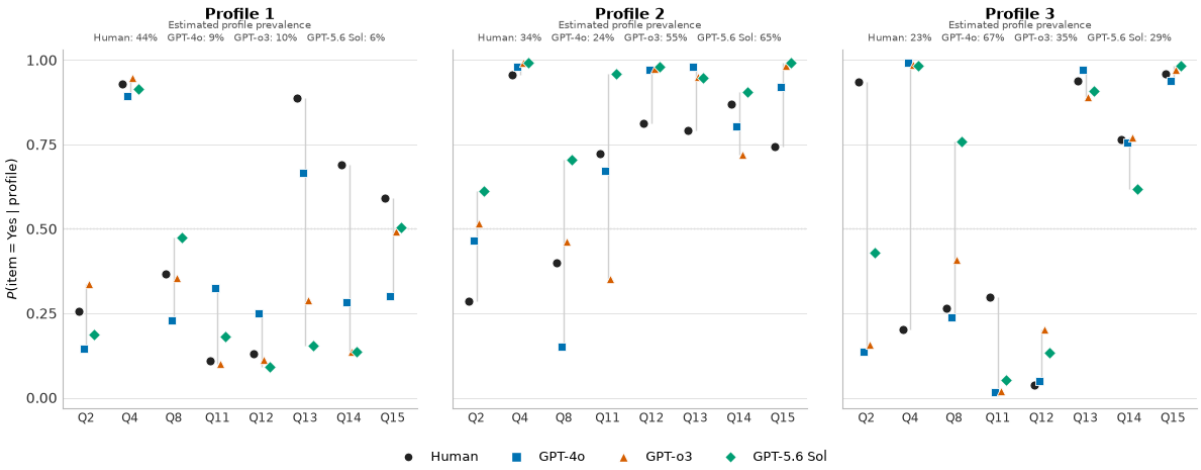


Figure 4: Profile-specific item-probability patterns for the human-derived and LLM-derived profiles after matching profile labels. Similar overall profile patterns do not necessarily imply similar transcript-level classifications.

profiles are used for assessment or feedback, future research should examine whether they are associated with meaningful differences in learner performance, patient outcomes, or the usefulness of feedback.

6 Limitations and Future Work

The study has several limitations. Each LLM was evaluated only once per transcript, so call-to-call variability was not assessed. The binary checklist also produced ceiling and floor effects that limited some items' usefulness for profile modeling. In addition, we used a relatively simple prompting approach; future work should evaluate more advanced prompt-engineering and calibration strategies. Finally, the data come from a single simulated respiratory dataset, so replication across specialties, populations, and real clinical settings is needed.

7 Conclusion

LLMs showed reasonable overall agreement with human ratings, but overall agreement alone can be misleading because substantial disagreement may remain on particular communication behaviors. Item-level analyses are therefore important for identifying which aspects of the rubric are evaluated similarly by human and LLM raters and which are not.

These differences also matter when item ratings are combined to identify communication profiles.

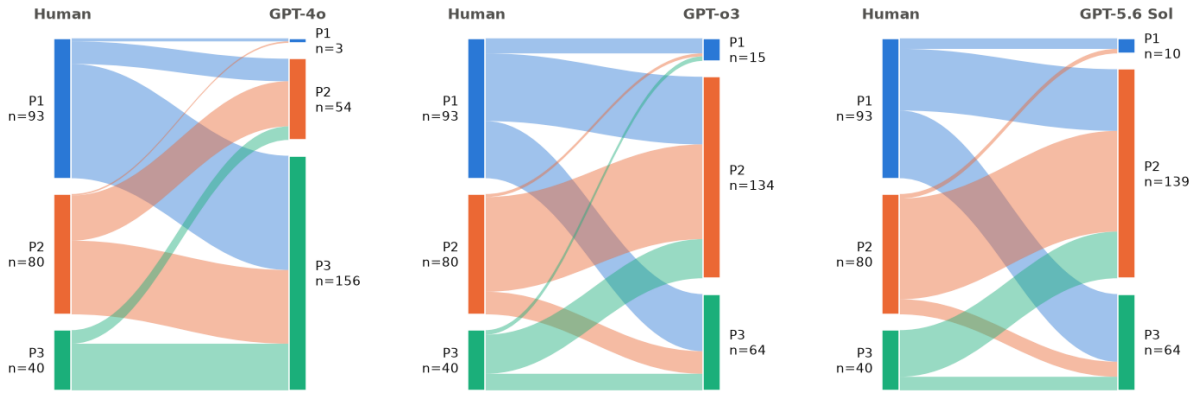


Figure 5: Human-to-LLM profile-classification transitions for GPT-4o (left), GPT-o3 (center), and GPT-5.6 Sol (right), $N = 213$. Ribbon width represents the number of transcripts.

Although the LLM-derived profiles showed similarity in their overall patterns, substantial differences remained for particular item probabilities, and the models often classified transcripts differently from the human-derived solution. Thus, when profile-based assessment or feedback is the intended use, evaluation should consider not only item-level agreement but also whether human and LLM ratings produce similar profile structures and transcript-level classifications.

Acknowledgments

We thank Faria Ali, Mudra Kale, and Sohail Shaikh for their contributions as human coders for this study.

References

Habibeh Ahmadipour and Nahid Mortazavi. 2015. Medical error and related factors during internship and residency. *Indian J Med Ethics*, 12(4):215–9.

Tavinder Ark, Adina Kalet, Linda Tewksbury, Lisa Altshuler, Ruth Crowe, Jeffrey Wilhite, Khemraj Hardowar, Sondra Zabar, and Colleen Gillespie. 2024. Validity evidence for the clinical communication skills assessment tool (ccsat) from 9 years of implementation in a high stakes medical student osce. *Patient Education and Counseling*, 127:108323.

Michael T. Brannick, H. Tugba Erol-Korkmaz, and Matthew Prewett. 2011. A systematic review of the reliability of objective structured clinical examination scores. *Medical Education*, 45(12):1181–1189.

Nilesh Chavda, Priti Solanky, Jatin V. Dhanani, Aashal Shah, Nirav Patel, and Sunilkumar Bhadiyadara. 2020. Assessment of clinical communication skills of medical students through the simulated patient approach. *Journal of Medical Education*, 19(3):e108661.

A. Angel Dalus, J. Gorshia, and M. P. Brundha. 2024. Enhancing physician-patient communication: The impact of the kalamazoo consensus statement on medical education. *International Journal of Medicine and Public Health*, 14(4):647–651.

Shih-Chen Fan, Shao-Tong Tsai, Yi-Ching Wang, Meng-Lin Lee, Sheau-Ling Huang, and Ching-Lin Hsieh. 2024. Reliability of the gap–kalamazoo communication skills assessment form in occupational therapy. *British Journal of Occupational Therapy*, 87(7):424–433.

Faiha Fareez, Tishya Parikh, Christopher Wavell, Saba Shahab, Meghan Chevalier, Scott Good, Isabella De Blasi, Rafik Rhouma, Christopher McMahon, Jean-Paul Lam, Thomas Lo, and Christopher W. Smith. 2022. A dataset of simulated patient-physician medical interviews with a focus on respiratory cases. *Scientific Data*, 9(1):313.

Jadon Geathers, Yann Hicke, Colleen Chan, Niroop Rajashekar, Sarah Young, Justin Sewell, Susannah Cornes, Rene F Kizilcec, and Dennis Shung. 2025. Benchmarking generative ai for scoring medical student interviews in objective structured clinical examinations (osces). In *International Conference on Artificial Intelligence in Education*, pages 231–245. Springer.

Kilem Li Gwet. 2008. Computing inter-rater reliability and its variance in the presence of high agreement. *British Journal of Mathematical and Statistical Psychology*, 61(1):29–48.

Matthew D. Hoffman and Andrew Gelman. 2014. The no-u-turn sampler: Adaptively setting path lengths in hamiltonian monte carlo. *Journal of Machine Learning Research*, 15(47):1593–1623.

Reyhaneh Hosseinpourkhoshkbari and Richard M. Golden. 2025. Clinical reasoning assessment methods: Current simulated environment practice and future prospects using ai and psychometrics. *SSRN Electronic Journal*.

Kelly V Johnson, Allison L Scott, and Lisa Franks. 2020. Impact of standardized patients on first semester nursing students self-confidence, satisfaction, and communication in a simulated clinical case. *SAGE open nursing*, 6:2377960820930153.

J. Richard Landis and Gary G. Koch. 1977. The measurement of observer agreement for categorical data. *Biometrics*, 33(1):159–174.

Yuelin Li, Jennifer Lord-Bessen, Mariya Shiyko, and Rebecca Loeb. 2018. Bayesian latent class analysis tutorial. *Multivariate Behavioral Research*, 53(3):430–451.

Gregory Makoul. 2001. Essential elements of communication in medical encounters: the kalamazoo consensus statement. *Academic medicine*, 76(4):390–393.

Brian Mavis, Jane Turner, Kathryn Lovell, and Dianne Wagner. 2006. Developments: faculty, students, and actors as standardized patients: expanding opportunities for performance assessment. *Teaching and learning in medicine*, 18(2):130–136.

HyeRin Roh, Kyung Hye Park, and Song Yi Park. 2017. Verbal communication of students with high patient–physician interaction scores in a clinical performance examination assessed by standardized patients. *Korean Journal of Medical Education*, 29(4):241–251.

John Salvatier, Thomas V. Wiecki, and Christopher Fonnesbeck. 2016. Probabilistic programming in python using pymc3. *PeerJ Computer Science*, 2:e55.

Ameer Hamza Shakur, Michael J Holcomb, David Hein, Shinyoung Kang, Thomas O Dalton, Krystle K Campbell, Daniel J Scott, and Andrew R Jamieson. 2024. Large language models for medical osce assessment: a novel approach to transcript analysis. *arXiv preprint arXiv:2410.12858*.

Hardeep Singh, Traber Davis Giardina, Ashley ND Meyer, Samuel N Forjuoh, Michael D Reis, and Eric J Thomas. 2013. Types and origins of diagnostic errors in primary care settings. *JAMA internal medicine*, 173(6):418–425.

Michelle Yoon and Veronica Michaelsen. 2015. Critical synthesis package: The kalamazoo consensus statement assessment tools. *MedEdPORTAL*, 11:10098.

A Annotation Prompt

LLM Communication-Rating Prompt

Role. You are an expert medical communication rater.

Your task is to read a clinical interview transcript and evaluate **only the doctor's communication behavior** using the coding rules below.

Important instructions

- Doctor utterances begin with `D:`.
- Patient utterances begin with `P:`.
- Score **only** the doctor's behavior.
- Use the patient's responses only as context.
- Do not infer behaviors that are not supported by the transcript.
- Return binary scores only:
  - 1 = YES, behavior is present
  - 0 = NO, behavior is absent
- Apply the coding rules exactly as written below.

Q1. Did the doctor greet the patient?
Score 1 if the doctor acknowledges or greets the patient before beginning health-related questions.
Examples:

- "Hello"
- "Hi"
- "How are you?"

Otherwise, score 0.

Q2. Did the doctor use words that show care and concern throughout the interview?
Score 1 if the doctor demonstrates compassion, empathy, care, or concern beyond routine medical questioning.
This may include responding compassionately when the patient discusses:

- A traumatic event
- Death of a loved one
- Emotional distress
- Fear or worry
- Other difficult personal experiences

Otherwise, score 0.

Q3. Did the doctor allow the patient to complete their opening statement without interruption?
Because the transcripts are text-based and interruptions cannot be reliably assessed, always score:

$$Q3 = 1$$

Q4. Did the doctor ask "Is there anything else?" or otherwise elicit a full set of concerns?
Score 1 if the doctor asks sufficiently broad and comprehensive background or follow-up questions to obtain a full picture of the patient's problem.
Examples include asking:

- What makes the symptom better or worse?
- Whether other symptoms occurred before or with the main symptom
- About associated symptoms
- About symptom progression
- Other relevant medical questions beyond the main symptom itself

For example, for a patient with a cough:

- "Does anything relieve the cough?"
- "Does anything make the cough worse?"
- "Did another symptom appear before the cough?"

Otherwise, score 0.

Q5. Did the doctor explain and/or negotiate an agenda for the visit?
Score 1 if, near the beginning of the interview, the doctor:

- Explains the purpose of the visit
- Explains what the patient is being seen for
- Describes what will be discussed
- Explains how the appointment may proceed

Otherwise, score 0.

Q6. Did the doctor begin with the patient's story using open-ended questions?
Score 1 if the doctor begins by inviting the patient to describe the problem or reason for the visit in their own words.
Examples:

- "What brings you in today?"
- "Tell me about your symptoms."
- "Can you tell me what has been going on?"

Otherwise, score 0.

Q7. Did the doctor clarify details as necessary using more specific questions?
Score 1 if, after the patient describes the main problem, the doctor asks more specific questions to clarify symptoms, history, risk factors, or related details.
Examples:

- "Is anyone around you sick?"
- "Have you been traveling?"
- "Have you lost or gained weight?"
- Questions about symptom characteristics or associated symptoms

Otherwise, score 0.

Q8. Did the doctor summarize information and give the patient an opportunity to correct or add information?
Score 1 if the doctor summarizes or restates information provided by the patient, particularly toward the end of the interview. The summary should reflect the patient's reason for the visit, symptoms, or information previously discussed.
Otherwise, score 0.

Q9. Did the doctor transition effectively to additional questions?
Score 1 if the doctor's questions follow a logical and organized sequence.
The doctor should generally finish exploring one symptom or topic before moving to another rather than jumping abruptly between unrelated topics.
Example: For a patient with headache, the doctor may first complete headache-related questions and then move to related vision or other neurological symptoms.
Otherwise, score 0.

Q10. Did the doctor ask about life events, circumstances, or other people that might affect the patient's health?
Score 1 if the doctor asks about relevant contextual factors that may affect the patient's health.
Examples:

- Travel
- Sick contacts
- Stress
- Family medical history
- Social circumstances
- Environmental exposures

- Other relevant life events

Otherwise, score 0.

Q11. Did the doctor elicit the patient's beliefs, concerns, and expectations about the illness and treatment?
Score 1 if the doctor identifies, explores, or meaningfully responds to the patient's beliefs, concerns, fears, expectations, or worries about the illness or treatment.
Example: If a patient says they are afraid to take a medication because it causes brain fog, score 1 if the doctor acknowledges and addresses that concern with explanation, discussion, or empathy.
Do not score 1 merely because the doctor provides medical information. There must be evidence that the patient's own belief, concern, expectation, or worry was addressed.
Otherwise, score 0.

Q12. Did the doctor respond explicitly to the patient's statements about ideas, feelings, or values?
Score 1 if the patient expresses an idea, belief, feeling, fear, or personal value and the doctor directly responds to it.
Examples of appropriate responses include:

- Explanation
- Reassurance
- Empathy
- Discussion of the patient's concern

Example: If the patient says they are worried their symptoms may indicate cancer, score 1 if the doctor directly responds to that concern with explanation and/or empathy.
Otherwise, score 0.

Q13. Did the doctor explain using words that are easy for the patient to understand?
Score 1 if the doctor uses clear, everyday, patient-friendly language when explaining:

- Symptoms
- Diagnoses
- Tests
- Procedures
- Treatment
- Next steps

Score 0 if the explanation relies mainly on unexplained medical jargon.
If the doctor does not provide any meaningful explanation, score 0.

Q14. Did the doctor ask whether the patient had any questions, concerns, or other issues?
Score 1 if, usually near the end of the interview, the doctor explicitly asks whether the patient:

- Has any questions
- Has any concerns
- Wants to add anything
- Has another issue to discuss
- Needs clarification

Otherwise, score 0.

Q15. Did the doctor acknowledge the patient and appropriately close the interview?
Score 1 if the doctor ends the interview appropriately rather than stopping abruptly.
An appropriate close may include:

- A conclusion
- Explanation of next steps
- Follow-up instructions
- Safety instructions
- A closing or acknowledgement of the patient

Otherwise, score 0.

LLM-Based Pairwise Judgment for Math Item Parameter Modeling using Workflows and Agents

Suhwa Han and Frank Rijmen

Cambium Assessment

Correspondence: suhwa.han@cambiumassessment.com

Abstract

This study examines the feasibility of using large language models (LLMs) as judges for pairwise comparisons of item difficulty and to the extent which the resulting comparison outcomes recover banked difficulty parameters. The study in particular fine-tunes an instruction-tuned LLM to evaluate the impact of task-specific fine-tuning on the parameter prediction accuracy. The study also demonstrates an agentic approach to hyperparameter tuning of LLM-training task.

1 Introduction

Item parameter modeling in educational measurement refers to the prediction of item response theory (IRT) parameters using item content available—such as item text—prior to operational administration (AlKhuzaey et al., 2024). With recent advances in large language models (LLMs), researchers have increasingly leveraged LLMs to develop models that can predict item parameters using item content. One prominent line of work involves fine-tuning of encoder-based LLMs with regression heads using banked item parameters as labels (Maeda and Lu, 2026; Han et al., 2025; Yaneva et al., 2024). Other approaches exploit the advanced capabilities of foundation models, such as GPT-4 (OpenAI, 2023), by prompting them to directly elicit numerical ratings of item difficulty (Razavi and Powers, 2025). While each of these approaches has reported promising results, both have limitations. The encoder-based fine-tuning approaches are generally limited by smaller model sizes (hundreds of millions of parameters) compared to generative LLMs with multi-billion parameters. For prompting-based approaches with foundational models, prior work has shown that scalar ratings elicited from LLMs are prone to biased or compressed score distributions (Licht et al., 2025; Stureborg et al., 2024).

More recently, interest has emerged in using LLMs for pairwise comparison methods (Han and Rijmen, 2026; Li and Gushta, 2026; Smith and Rewley, 2026). Rather than eliciting absolute difficulty ratings from LLMs, these methods require only a series of relative judgments—which item in a pair is more difficult—a task generally considered more straightforward (Thurstone, 2017). The resulting comparison outcomes are then aggregated using a measurement model such as the Bradley-Terry model (Bradley and Terry, 1952) to recover a continuous latent difficulty scale. While pairwise comparison methods with human judges have previously been applied for item parameter prediction (Australian Council for Educational Research, 2022), applying these methods to large item pools with human judges remain practically challenging due to the quadratically growing number of required judgments. LLM-based pairwise judgments offer a scalable solution by automating the repetitive comparisons. Moreover, LLM-based pairwise methods may address aforementioned limitations of the existing LLM-based methods. First, they are well suited to large generative models because pairwise comparisons can be naturally formulated as next-token prediction tasks. Second, they can bypass the problem of compressed or biased score distributions associated with direct scalar ratings from LLMs as the latent scale is instead derived through a measurement model.

1.1 Study Purpose and Contributions

The purpose of the current study is to examine the feasibility of using LLMs as judges to automate pairwise comparison tasks and the extent to which item parameters can be recovered by aggregating LLM-generated pairwise outcomes using the Bradley-Terry model. Specifically, we fine-tune an open-source generative LLM for the pairwise comparison task using a dedicated training set with which the true pairwise outcomes from banked item

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 248–259
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

parameters serve as training labels. We then evaluate the fine-tuned model's performance on the pairwise task by applying it to a separate development set and assessing the closeness of the latent scale derived from pairwise outcomes to the banked parameter values. Additionally, we compare the performance of the fine-tuned model to that of the off-the-shelf counterpart applied to the same development set via inference only (i.e., no task-specific fine-tuning).

To the authors' knowledge, no prior study has fine-tuned an open-source generative LLM to develop LLM judges for the item parameter modeling. This distinction is particularly relevant for this pairwise comparison methods where computational overhead grows quadratically with the number of items to be compared. While foundational models could be theoretically applied, doing so at scale would result in substantial token costs, rendering the fine-tuning open-source models a more scalable alternative.

A further contribution of this work is the explicit characterization of the research pipeline as an *agentic system*. Following the taxomony of (Anthropic, 2024), we distinguish *workflows*—in which LLMs are orchestrated through predefined paths—and *agents*—in which LLMs dynamically direct their own processes and tool usage—and describe relevant components of the pipeline according to this distinction. We intend this characterization to serve as early guidance for the systematic and efficient use of agentic systems in LLM-based research avenue in the field.

2 Using LLMs as Pairwise Judges for Item Parameter Modeling

Substituting human judges with LLM judges warrants special consideration. This section discusses two such considerations that informed the design of the current study's methodology on item parameter modeling.

2.1 Positional Bias in LLMs as Pairwise Judges and the Swap Test

Prior studies on LLMs as pairwise judges have repeatedly reported systematic positional bias, whereby the model's judgment is swayed by the order in which items are presented rather than by their substantive content (Wang et al., 2024; Shi et al., 2025). In the context of item parameter modeling, Han and Rijmen (2026) showed that this positional bias persisted even when a foundational model was used for the comparison, and that it can be mitigated by prompting the model to generate explanatory texts before the verdict (Wang et al., 2024). As a metric to quantify a model's positional bias, Shi et al. (2025) proposed a swap test in which every pair is evaluated twice with item presentation order reversed between passes. A comparison passes the swap test if the model selects the same item regardless of presentation order while inconsistent decisions indicate a positionally biased judgment.

2.2 Generative LLMs in Pairwise Judgments

Generative LLMs are trained autoregressively, meaning each token in a sequence is predicted conditioned on all preceding tokens. This unidirectional architecture naturally makes these models well-suited for text generation tasks, which can be elicited by wrapping inputs in structured, chat-formatted prompts. Pairwise comparison tasks, which require only a single discrete output—a token indicating which item is more difficult (e.g., "A" or "B")—can be naturally formulated as a next-token prediction task without any architectural modification required. At inference, the model's judgment is obtained by selecting the token with the highest probability at the designated output position, if greedy decoding is chosen. Fine-tuning can be formulated in the same manner by using the true pairwise outcome as the training target for that same token position.

3 Methods

This section details the methods used in this study's item parameter modeling pipeline, from data preparation, prompt design, model training and inference to Bradley-Terry aggregation of the resulting pairwise outcomes.

3.1 Data and Sampling

This study utilized 1,247 of 1,386 Grade 6 mathematics items drawn from a large-scale K-12 assessment. Each item was associated with banked IRT parameters calibrated from operational administration. The 1,247 items were partitioned into three subsets: a training set (975 items), an inner validation set used to select the best checkpoint during fine-tuning (133 items), and a development set used for model selection and for reporting the results (139 items). The remaining 139 items constitute a test partition that was reserved and not evaluated

in the present study (see limitations in Section 6). The majority of items were dichotomous (1,213 items) with a small subset scored polytomously (34 items). Dichotomous items were calibrated using two-parameter logistic (2PL) model (Birnbaum, 1968) while polytomous items were calibrated with generalized partial credit model (GPCM) (Muraki, 1992). Appendix A shows the descriptive statistics of the banked difficulty parameters.

Evaluation Target: Banked IRT Difficulty. The evaluation target in this study was the banked IRT difficulty parameter. Because the 2PL model is a special case of GPCM, 2PL-calibrated parameters were treated within the same GPCM parameterization in this study. GPCM describes the probability of an examinee with a latent trait level θ to obtain a score of $v \in \{0, 1, \ldots, m_i\}$ for item i as:

$$p_{iv} = \frac{\exp\left(\sum_{r=0}^{v} Da_i\left(\theta - b_i + d_{ir}\right)\right)}{\sum_{c=0}^{m_i} \exp\left(\sum_{r=0}^{c} Da_i\left(\theta - b_i + d_{ir}\right)\right)},$$

where a_i and b_i respectively denote the overall discrimination and difficulty parameters for item i, and d_{ir} represents the step parameter for the category r for the item. The GPCM parameters used in this study were estimated under the constraints $d_{i0} = 0$ and $\sum_r d_{ir} = 0$. For dichtomous items, the model reduces to 2PL. Accordingly, the item-level difficulty parameter b_i was used as the evaluation target for both dichotomous and polytomous items.

Augmented Item Text. As the input to feed to the LLM prompt for the pairwise comparison task, we constructed a single augmented item text per item by concatenating three distinct text sources: (i) *item text*, the raw text of the question and stimulus presented to students, (ii) *accessibility context*, additional item context from accommodation features that provides information often not captured in item text alone (e.g., text descriptions of images for Braille users), and (iii) *scoring rubric* that contains the answer key or the human-readable scoring rubric. Appendix B shows an illustrative example of this augmented item text for a synthetic item.

3.2 Language Model

Throughout this study, we used Gemma-4-12B (Gemma Team, 2026), an open-source, instruction-tuned generative language model with approximately 12 billion parameters. As an instruction-tuned model, the model takes inputs formatted as a structured chat conversation consisting of system, user and assistant turns. We used this chat format to structure every generation step in the pipeline.

3.3 LLM Difficulty Breakdown for Bias-free Pairwise Comparison

Motivated by the finding of Wang et al. (2024) that positional bias can be mitigated by eliciting explanations from LLMs before a verdict, we adopted a variant of chain-of-thought (CoT) prompting (Zhou et al., 2022) where the model produces a structured *difficulty breakdown* for each item prior to comparison. The breakdown is generated through a sequence of prompts, each addressing a distinct perspective on the factors that may contribute to item difficulty. The final breakdown contains a sequence of the model's responses from each step, incorporating the responses of preceding prompts as context. This breakdown structure not only ensures that explanatory reasoning precedes the verdict, but may also yield a more comprehensive difficulty breakdown based on multiple perspectives.

Specifically, our CoT prompting strategy consisted of three stages: (i) *expert perspective* where the model takes the perspective of a math content expert evaluating the item's content and cognitive demand, (ii) *student perspective* where the model takes the perspective of a typical student and examines the required prerequisite and the effect of response format, and (iii) *synthesis* where the model integrates the preceding responses into an overall difficulty level and drivers. Appendix C and D respectively show examples of our CoT prompts for a synthetic item and the resulting final breakdown.

3.4 Pairwise Judgment, Swap Test, and Tie-breaking

Once structured breakdowns had been generated for individual items, pairwise judgments were elicited by presenting the model with the breakdowns of two items and asking it to select the more difficult item based on the breakdowns. Each pair was judged twice with the presentation order reversed. A swap test was then applied to assess the model's predisposition to positional bias, checking whether the model selected the same item as more difficult across both passes despite the reversed order. When the two passes disagreed, the pair was flagged as positionally biased and resolved through a dedicated tie-breaking prompt. Appendix E and F respectively show examples of the pairwise comparison prompt and the tie-breaking prompt used when a pair fails the swap test.

3.5 Estimating Latent Difficulty via Bradley-Terry Model

Once pairwise judgments had been obtained for all evaluated item pairs, the Bradley-Terry model (Bradley and Terry, 1952) was used to aggregate the outcomes to place the items on a continuous latent difficulty scale. The model assumes that each item i has a latent strength parameter β_i and describes the probability that item i is judged more difficult than item j as

$$P(i > j) = \frac{\exp(\beta_i)}{\exp(\beta_i) + \exp(\beta_j)}.$$

The model parameters were estimated using the `choix` Python package, yielding each item's position on a continuous latent difficulty scale. Because the Bradley-Terry estimates are identified only up to an additive constant, the parameters were linearly rescaled to match the mean and standard deviation of the Bradley-Terry estimates to those of the banked difficulty parameters.

3.6 Fine-tuning a Generative Model for Pairwise Comparison

The difficulty breakdown, pairwise judgment and the Bradley-Terry latent difficulty estimation described in Sections 3.3 to 3.5 together constitute the *inference* pipeline. For training, we targeted a specific component of the pipeline—the pairwise judgment step in 3.4—as the trainable signal in scaling the difficulty parameters essentially come from the pairwise comparison task.

Training Objective. The training examples from the designated training data were prepared using the same chat format used at inference (see Appendix E): a user turn containing the pairwise comparison prompt followed by an assistant turn containing the target response. For training, the verdict token in the target response was derived from the banked difficulty parameters. The assistant turn, on the other hand, included a fixed, templated sentence identical across all examples. The loss was computed over the tokens in the entire assistant turn using the standard next-token cross-entropy objective.

Training Data Sampling. One key consideration in training on a pairwise task is the number of training samples. Because the label is defined at the pair level, the number of possible training pairs grows quadratically, reaching as high as $\binom{n}{2}$ pairs for n items. With 975 training items, the size of training samples can amount to 474,825 pairs. Training on these many samples is not only unnecessary but also inefficient as each item recurs across $n - 1$ pairs and incurs highly redundant information. To address this, we adopted a structured sampling approach that curates a subset.

Our structured sampling was conducted in three steps: (i) *margin filtering*, (ii) *stratification by difficulty gap* and (iii) *stub matching sampling*. The margin filtering considers the idea that a pair whose difficulty difference falls within the estimation error of the banked parameters provides noise rather than signal. For each candidate pair (i, j), we thus computed a standardized gap $d_{ij} = |b_i - b_j| / \sqrt{\mathrm{se}_i^2 + \mathrm{se}_j^2}$, and retained only pairs with $d_{ij} \geq 1$. Next, we stratified the remaining pairs by their difficulty gap $|b_i - b_j|$ into ten strata using equal- percentile boundaries. Finally, within this stratified structure, we sampled pairs such that each item appeared approximately $K = 80$ times in the training set. This sampling was done through a stub-matching procedure (Newman et al., 2001), where each item is assigned K stubs distributed across the strata and stubs are randomly paired with other stubs within the stratum.

Training Configuration. To fine-tune the Gemma model, we used Low-Rank Adaptation (LoRA; Hu et al., 2021) with a rank of 16 and a scaling factor of 32, applying adapters to the attention and projection matrices. We used a batch size of 1 with gradients accumulated over 16 steps. To guard against overfitting while selecting the best-performing model, we examined pairwise prediction accuracy—specifically accuracy on the verdict token—on the validation set of 133 items every 100 optimization steps. Training was terminated early if validation accuracy did not improve over 10 consecutive evaluations, and the final model was selected from the checkpoint with the highest accuracy as long as it shows decent diagnostics (see Figure 1 for details). In addition to early stopping, training was capped at 1 epoch as our preliminary analysis showed that completing a full epoch was unnecessary and led to degraded performance. Given that individual items reappeared many times across the training set in pairwise comparison, it is reasonable that optimal performance was achieved before a full epoch.

Learning Rate Optimization Using Agents. All the hyperparameters above were fixed a pri-

ori, while the *learning rate*, which arguably is one of the most important hyperparmeter, was selected through a dedicated search procedure using an agent, which will be detailed in Section 4.2.

4 Workflows and Agents in the Research Pipeline

We now characterize the current study through a lens of agentic system, following the framework of Anthropic (2024). The pipeline implemented here requires a continuous stream of LLM calls. When a task is complex enough to warrant multiple LLM calls and yet its steps can be specified deterministically, a *workflow* is appropriate. An *agent* is reserved for open-ended tasks whose steps depend on decisions that cannot be specified in advance and the LLM must dynamically determine how to proceed based on information obtained during execution. Figure 1 depicts how both workflows and an agent are incorporated into the current research pipeline.

4.1 Workflows in the Pipeline

As shown in Figure 1, most of the LLM components in the current pipeline constitute a workflow: the pipeline proceeds as a linear chain where each LLM call processes the output of the preceding step along a deterministic path. For example, the pairwise judgment operates on the difficulty breakdowns produced from the preceding LLM call while the tie-breaker gets invoked conditioning on the preceding swap test results. Importantly, however, this branching decision is entirely governed by the prescribed script rather than by the LLM itself. That is, although the LLM outputs provide information required by subsequent steps, the model itself never determines which action is executed next.

4.2 An Agent for Learning Rate Optimization

An *agent*, in contrast, dynamically determines its course of action. Learning rate optimization provides a natural setting for such an agentic approach. Selecting an appropriate learning rate is a critical component of LLM fine-tuning and yet such hyperparameter tuning has traditionally relied on manual experimentation (Bischl et al., 2023; Chaudhry and Mizuno, 2024). Although grid search offers a more systematic alternative, it can become computationally prohibitive for large generative models and may expend unnecessary resources for evaluating candidates that are unlikely to perform well. An agent can offer an adaptive solution. Rather than exhaustively sweeping a predefined set of candidates or relying on repeated human intervention, the agent can propose the next learning rate by intelligently interpreting evidence from previous runs, such as validation performance and training diagnostics.

Figure 1 depicts the learning rate search procedure in which the agent dynamically directs each successive learning rate. Operating on a specification written in a natural language (e.g., a markdown file) that defines its goal, constraints, and guidance on what evidence to weigh, the agent proceeds through a dynamic loop until it determines that sufficient experiments have been conducted to identify an optimal learning rate. Specifically, at each iteration, the agent invokes a *sub-workflow* that (i) conducts a smoke test to verify that the verdict token generation is done properly and (ii) performs the Bradley-Terry estimation on the development set to provide information about predictive performance. The agent then evaluates the outputs of this sub-workflow together with training diagnostics, such as gradient norms, to assess training stability and decides the next learning rate to test. For example, if substantial spikes in gradient norms are observed, the agent may suggest a lower learning rate.

5 Results

We next present the model performance results. Model performance was evaluated by examining how closely the latent scale estimated from the model's pairwise outcomes through the Bradley-Terry estimation reproduced the banked difficulty parameters. The Spearman and Pearson correlations and the root mean squared error (RMSE) were computed to guage this. In addition to prediction accuracy, we report swap-test pass rates (i.e., the proportion of item pairs for which the model consistently selected the same item despite the reversed presentation order) as a metric to assess the model's robustness to positional bias. All results reported below were computed on the development set (except for verdict-token accuracy computed on the inner validation set).

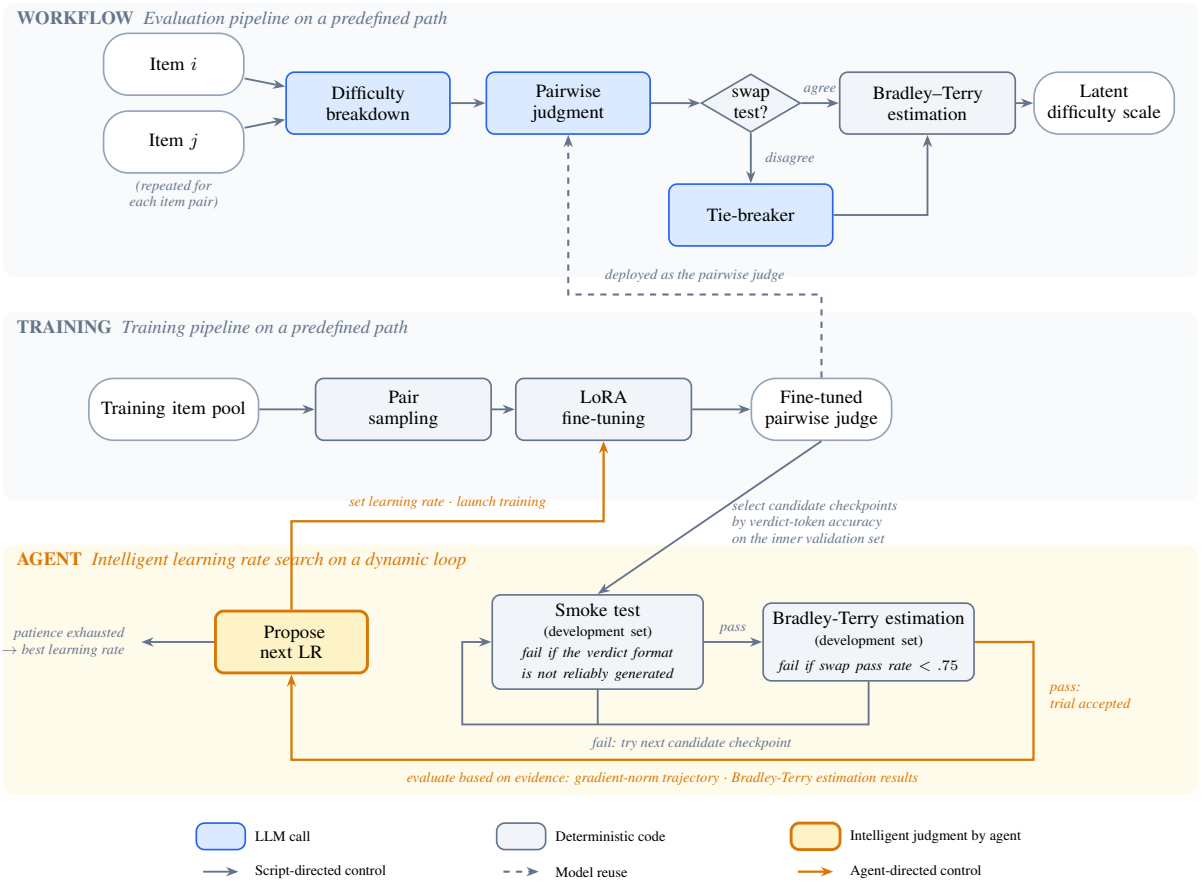


Figure 1: The research pipeline described as an agentic system

5.1 Difficulty Prediction Across the Learning Rate Search by Agent

Table 1 summarizes the difficulty prediction results across the learning rates proposed by the agent. The results show that the agent successively halved the learning rate, despite an apparently acceptable prediction outcome at Trial 1 (i.e., $\rho = .702$). This decision was driven by the training diagnostics: extreme spikes in the gradient norm and evaluation loss together with the failure of two candidate checkpoints on the smoke test led the agent to conclude that training was unstable at that learning rate (see Appendix G for datails). As the learning rate was halved, the correlations improved monotonically, from .702 to .766 in Spearman correlation and from .661 to .730 in Pearson correlation. Verdict-token accuracy increased monotonically as well although the gain from Trial 2 to 3 was marginal. Swap-test pass rates were uniformly high across all trials, indicating that the fine-tuned models were robust to positional bias.

Trial	Learning rate	ρ	r	Swap pass	Verdict-token accuracy
1	1×10^{-4}	.702	.661	.922	.813
2	5×10^{-5}	.747	.709	.914	.824
3	2.5×10^{-5}	.766	.730	.925	.826

Table 1: Difficulty prediction results across the learning rate search by the agent

5.2 Difficulty Prediction by Model

Table 2 presents difficulty prediction results for three models. In addition to the models based on the pairwise comparison method, results based on a fine-tuned DeBERTa model following the method of Han et al. (2025) are presented for comparison. The fine-tuned pairwise model predicted the banked difficulty parameters most closely ($\rho = .766$, $r = .730$, RMSE $= 1.057$), with the encoder fine-tuning method slightly behind on all reported metrics. By contrast, the effect of task-specific fine-tuning within the pairwise comparison framework was pronounced. Applying the identical evaluation pipeline with an off-the-shelf model achieved only $\rho = .448$ and $r = .401$. The higher swap-test pass rate of the fine-tuned model (.925 vs. .883)

further suggests that fine-tuning also improved the positional consistency of the generative model.

Model	ρ	r	RMSE	Swap pass
Pairwise, fine-tuned	.766	.730	1.057	.925
Pairwise, off-the-shelf	.448	.401	1.573	.883
Encoder fine-tuning	.733	.693	1.127	—

Table 2: Difficulty prediction results by model

6 Discussion and Limitations

This study examined the feasibility of using LLMs as judges for pairwise comparison of item difficulty and the extent to which the resulting outcomes can predict banked difficulty parameters when aggregated through the Bradley-Terry model. Using operational Grade 6 mathematics items from a large-scale K–12 summative assessment, the results showed that banked difficulty parameters can be recovered with adequate accuracy when the LLM was appropriately fine-tuned for the pairwise comparison task. In addition, the study characterized the research pipeline through the lens of an agentic system and demonstrated how the learning-rate selection task can be delegated to an agent. Guided by a specification document written in a natural language, the learning-rate-search agent autonomously selected successive candidate values by interpreting training dynamics and downstream prediction performance. This approach provides an adaptive alternative to manual tuning or exhaustive grid search, with the potential to improve both the efficiency of hyperparameter optimization and the use of computational resources in LLM fine-tuning.

This study is not without limitations. First, the current study remains in progress, and the fine-tuned model has not yet been evaluated on a completely held-out test set. Evaluation on the held-out set will provide a more rigorous assessment of generalization and help determine whether the learning-rate exploration by agent translates to improved performance beyond the development data. Second, the procedure evaluated in this study compares all possible item pairs within a pool, making it most directly applicable to placing an existing set of items on a common scale. For the more common scenario of placing a smaller number of new items relative to an existing calibrated pool, an anchoring design could instead be used where new items are compared against items with known difficulty values and the anchor-item parameters are fixed at their known values during the model estimation.

Disclaimer

Generative AI tools, including Claude and Gemini, were used to assist with writing and coding during the development of this study. All AI-assisted content and code were reviewed and validated by the authors, who retain full responsibility for the final work.

References

Samah AlKhuzaey, Floriana Grasso, Terry R Payne, and Valentina Tamma. 2024. Text-based question difficulty prediction: A systematic review of automatic approaches. *International Journal of Artificial Intelligence in Education*, 34(3):862–914.

Anthropic. 2024. Building effective agents. https://www.anthropic.com/engineering/building-effective-agents. Anthropic Engineering Blog. Accessed: 2026-09-04.

Australian Council for Educational Research. 2022. Pairwise comparison method: Concept note. Technical report, Australian Council for Educational Research.

Allan Birnbaum. 1968. Some latent trait models. *Statistical theories of mental test scores*.

Bernd Bischl, Martin Binder, Michel Lang, Tobias Pielok, Jakob Richter, Stefan Coors, Janek Thomas, Theresa Ullmann, Marc Becker, Anne-Laure Boulesteix, and 1 others. 2023. Hyperparameter optimization: Foundations, algorithms, best practices, and open challenges. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 13(2):e1484.

Ralph Allan Bradley and Milton E Terry. 1952. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345.

Zan Chaudhry and Naoko Mizuno. 2024. Exptest: Automating learning rate searching and tuning with insights from linearized neural networks. *arXiv preprint arXiv:2411.16975*.

Gemma Team. 2026. Gemma 4 technical report. *arXiv preprint arXiv:2607.02770*.

Suhwa Han and Frank Rijmen. 2026. Scaling item difficulty using llm-based pairwise judgments: An application of the bradley-terry-luce model. Paper presented at the annual meeting of the National Council on Measurement in Education April 2026, Los Angeles, CA.

Suhwa Han, Frank Rijmen, Allison Ames Boykin, and Susan Lottridge. 2025. Leveraging fine-tuned large

language models in item parameter prediction. In *Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con): Full Papers*, pages 250–264.

Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.

Lanrong Li and Matthew Gushta. 2026. Using ai judges' paired comparisons to estimate item difficulty. Paper presented at the annual meeting of the National Council on Measurement in Education April 2026, Los Angeles, CA.

Hauke Licht, Rupak Sarkar, Patrick Y Wu, Pranav Goel, Niklas Stoehr, Elliott Ash, and Alexander Miserlis Hoyle. 2025. Measuring scalar constructs in social science with llms. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 32132–32159.

Hotaka Maeda and Yikai "EK" Lu. 2026. Multimodal test item parameter prediction from text, images, and metadata: Fusing together ai vision and language models. *Educational and Psychological Measurement*, 0(0):00131644261460779.

Eiji Muraki. 1992. A generalized partial credit model: Application of an em algorithm. *Applied psychological measurement*, 16(2):159–176.

Mark EJ Newman, Steven H Strogatz, and Duncan J Watts. 2001. Random graphs with arbitrary degree distributions and their applications. *Physical review E*, 64(2):026118.

OpenAI. 2023. GPT-4 technical report. *Preprint*, arXiv:2303.08774.

Pooya Razavi and Sonya J Powers. 2025. Estimating item difficulty using large language models and tree-based machine learning algorithms. *arXiv preprint arXiv:2504.08804*.

Lin Shi, Chiyu Ma, Wenhua Liang, Xingjian Diao, Weicheng Ma, and Soroush Vosoughi. 2025. Judging the judges: A systematic study of position bias in llm-as-a-judge. In *Proceedings of the 14th International Joint Conference on Natural Language Processing and the 4th Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics*, pages 292–314.

Nate Smith and Kelly Rewley. 2026. Predicting item statistics: Thurstone scaling with large language models. Paper presented at the annual meeting of the National Council on Measurement in Education April 2026, Los Angeles, CA.

Rickard Stureborg, Dimitris Alikaniotis, and Yoshi Suhara. 2024. Large language models are inconsistent and biased evaluators. *arXiv preprint arXiv:2405.01724*.

Louis L Thurstone. 2017. A law of comparative judgment. In *Scaling*, pages 81–92. Routledge.

Peiyi Wang, Lei Li, Liang Chen, Zefan Cai, Dawei Zhu, Binghuai Lin, Yunbo Cao, Lingpeng Kong, Qi Liu, Tianyu Liu, and 1 others. 2024. Large language models are not fair evaluators. In *Proceedings of the 62nd annual meeting of the association for computational linguistics (volume 1: Long papers)*, pages 9440–9450.

Victoria Yaneva, Kai North, Peter Baldwin, Le An Ha, Saed Rezayi, Yiyun Zhou, Sagnik Ray Choudhury, Polina Harik, and Brian Clauser. 2024. Findings from the first shared task on automated prediction of difficulty and response time for multiple-choice questions. In *Proceedings of the 19th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2024)*, pages 470–482, Mexico City, Mexico. Association for Computational Linguistics.

Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc Le, and 1 others. 2022. Least-to-most prompting enables complex reasoning in large language models. *arXiv preprint arXiv:2205.10625*.

A Descriptive Statistics of the Banked Difficulty Parameters

	Train ($n = 975$)	**Inner validation** ($n = 133$)	**Development** ($n = 139$)
Min	-3.934	-2.396	-1.900
1st Qu.	0.105	-0.125	-0.065
Median	0.921	0.853	0.678
Mean	0.933	0.944	0.856
3rd Qu.	1.720	1.772	1.637
Max	7.522	9.158	6.778
SD	1.400	1.593	1.442

Table 3: Descriptive statistics for the banked item difficulty parameter (b_i) across the subsets

B Example of the Augmented Item Text Input (Synthetic Item)

Augmented item text input

[Item Text: the question or prompt as presented to the student]

A bag contains red and blue marbles in a ratio of 3 to 5. If there are 25 blue marbles, how many red marbles are in the bag?

[Accessibility Context: spoken and Braille descriptions of images, graphs, and interactive elements not visible in the item text]

A double number line is shown. The top line is labeled “Red marbles” with tick marks at 0, 3, 6, 9, 12, and 15. The bottom line is labeled “Blue marbles” with tick marks at 0, 5, 10, 15, 20, and 25. Corresponding values are vertically aligned.

[Scoring Rubric: the answer key for simple items, or a full scoring rubric describing acceptable responses for complex item types]

Correct answer: 15 . 1 point for the correct value; 0 points otherwise. Equivalent unsimplified forms are not accepted.

Figure 2: Illustrative example of the augmented text input on a synthetic item

C Example of the Difficulty Breakdown Prompt Chain (Synthetic Item)

Difficulty breakdown prompt chain

Stage 1 — Expert perspective

User: Analyze the following math item:
{Item Text}

You are a math content expert. Analyze this item from an expert perspective:
(1) **Content Domain**: what concept(s) does this assess, and is it easy or hard for the intended students? (2) **Cognitive/Competency Demand**: how much structured reasoning, coordination of steps, or abstraction is required? (3) **Representational/Processing Demand**: how much translation between representations (notation, graphs, tables) is required? Be concise; use bullet points; max 120 words; do NOT solve the problem.

Assistant: - Content Domain: Ratios and proportional reasoning (Grade 6); on-grade, moderately familiar. - Cognitive Demand: Set up a part-to-part ratio and scale it to a given quantity; a few coordinated steps. - Representational Demand: Must read a double number line and relate the visual to the numeric ratio.

Stage 2 — Student perspective

User: Based on your previous analysis:
{previous_answers}

Now shift perspective. Imagine a typical student encountering this item:
(1) **Prior/Prerequisite Knowledge**: is the required knowledge common, emerging, or fragile for a typical student? (2) **Response Format**: does the format (multiple choice vs. open-ended) add or remove difficulty? Be concise; max 100 words.

Assistant: - Prior Knowledge: Equivalent-ratio reasoning is emerging for many Grade 6 students; scaling a ratio can be fragile. - Response Format: Open-ended numeric entry removes answer-elimination cues, adding difficulty relative to multiple choice.

Stage 3 — Synthesis

User: Based on all previous analysis:
{previous_answers}

Synthesize the overall difficulty for a typical student. State: (1) difficulty level (LOW / MODERATE / HIGH), (2) the 2 most significant difficulty drivers. Max 80 words.

Assistant: Difficulty level: MODERATE. Most significant drivers: (1) scaling a part-to-part ratio to a target quantity; (2) integrating the double-number-line representation with the numeric ratio.

Figure 3: Illustrative example of the three-stage difficulty breakdown prompt chain on a synthetic item

D Example of the Final Difficulty Breakdown (Synthetic Item)

Final difficulty breakdown

- Content Domain: Ratios and proportional reasoning (Grade 6); on-grade, moderately familiar.
- Cognitive Demand: Set up a part-to-part ratio and scale it to a given quantity; a few coordinated steps.
- Representational Demand: Must read a double number line and relate the visual to the numeric ratio.

- Prior Knowledge: Equivalent-ratio reasoning is emerging for many Grade 6 students; scaling a ratio can be fragile.
- Response Format: Open-ended numeric entry removes answer-elimination cues, adding difficulty relative to multiple choice.

Difficulty level: MODERATE.
Most significant drivers: (1) scaling a part-to-part ratio to a target quantity; (2) integrating the double-number-line representation with the numeric ratio.

Figure 4: The final difficulty breakdown on a synthetic item

E Example of the Pairwise Comparison Prompt

Pairwise comparison prompt

System: You are a precise and detail-oriented educational assessment expert.

User: You are comparing the difficulty of two math items based on detailed analyses.

ANALYSIS A: {difficulty breakdown of Item A}

ANALYSIS B: {difficulty breakdown of Item B}

INSTRUCTIONS:
- Based ONLY on the analyses above, determine which item (A or B) is MORE DIFFICULT for typical students
- Consider ALL difficulty factors mentioned in the analyses
- The order of presentation (A vs B) is completely arbitrary and should NOT influence your decision
- You MUST choose either A or B

RESPONSE FORMAT:
FINAL ANSWER: <Your choice of "A" or "B">
JUSTIFICATION: <1-2 sentences explaining the key difficulty factors>

Assistant: **FINAL ANSWER: B**
JUSTIFICATION: Both items are moderate, but Item B additionally requires multiplying two fractions and interpreting the product as an area, which is more error-prone than scaling a single part-to-part ratio.

Figure 5: Illustrative example of the pairwise comparison prompt

F Example of the Tie-Breaking Prompt

Tie-breaking prompt

System: You are a precise and detail-oriented educational assessment expert.

User: You are resolving a disagreement about which of two math items is more difficult.

Two previous comparisons of the same pair produced OPPOSITE conclusions:
- Comparison 1 conclusion: Item A is more difficult, owing to integrating the double-number-line representation with the numeric ratio.
- Comparison 2 conclusion: Item B is more difficult, because multiplying two fractions is more error-prone.

A: {difficulty breakdown of Item A}

B: {difficulty breakdown of Item B}

INSTRUCTIONS:
- The two previous comparisons disagreed. Look MORE CAREFULLY to break the tie.
- Focus on the factors most predictive of student difficulty: response format (multiple choice vs open-ended), prior-knowledge fragility, and cognitive demand.
- You MUST choose either A or B. Do NOT say they are equal.

RESPONSE FORMAT:
FINAL ANSWER: <Your choice of "A" or "B">
JUSTIFICATION: <1-2 sentences explaining what tipped the balance>

Assistant: **FINAL ANSWER: B**
JUSTIFICATION: Both items share open-ended formats and emerging prerequisites, but Item B's two-step fraction multiplication imposes a higher cognitive demand than scaling a single ratio.

Figure 6: Illustrative example of the tie-breaking prompt

G Training Diagnostics from the First Learning Rate Trial (1×10^{-4})

Step	Training loss	Maximum gradient norm	Evaluation loss	Verdict-token accuracy
100	2.308	114.52	.022	.721
200	0.024	0.82	.019	.769
300	0.021	0.99	.018	.819†
400	0.095	14.59	.020	.733
500	0.018	4.94	.020	.770
600	0.017	1.56	.019	.804
700	0.069	66.69	.021	.758
800	0.128	408.75	.823	.468
900	0.045	226.98	.017	.806
1000	0.014	3.88	.021	.786
1100	0.013	0.56	.021	.813†
1200	0.011	0.41	.019	.818†
1300	0.009	0.69	.019	.813

Table 4: Diagnostics from the first learning-rate trial (1×10^{-4}) as available to the agent. Training loss was the mean and the gradient norm the maximum over the 100 optimization steps preceding each evaluation. Early stopping terminated the run at step 1300 of a 2071-step ceiling. † marks the three checkpoints with the highest verdict-token accuracy, which the agent screened in descending order. Those at steps 300 and 1200 failed the smoke test, and the checkpoint at step 1100 was therefore selected as the final model for this trial.

Beyond Volume: How Cognitive Fingerprints Predict Residual Gain in AI Tutoring

Aleena K Raj
Learning Analytics & Psychometrics
Kaplan North America
aleena.kraj@kaplan.edu

Nina Deng
Learning Analytics & Psychometrics
Kaplan North America
nina.deng@kaplan.com

Abstract

This study evaluates an AI tutoring tool in a pre-licensure exam preparation product. Using residual gain modeling, elastic-net feature selection, and Gaussian Mixture clustering on LLM-derived cognitive fingerprints, the study identifies meaningful learner profiles and shows that the cognitive depth of student–AI interaction, not volume, drives measurable learning gain.

1 Introduction

Generative AI-based tutoring tools are increasingly integrated into online learning platforms. While conversational interfaces can enhance engagement and personalization, their efficacy remains underevaluated. Early research on intelligent tutoring systems often fails to isolate students' baseline capabilities from treatment effects (Shadish et al., 2002; Pelánek et al., 2016). Recent studies report highly heterogeneous results (Kestin et al., 2024; Mollick and Mollick, 2024), and some reveal that unguided generative AI use leads to cognitive offloading and degraded critical thinking (Bastani et al., 2025).

One challenge is that many evaluations report aggregate treatment effects, implicitly assuming that all students benefit uniformly from AI tutoring. When the intervention's effect depends fundamentally on *how* students use the tool—not merely *whether* they use it—aggregate metrics will systematically underestimate (or entirely miss) the tool's value for heterogeneous learners.

Motivated by this gap, the present study aims to address the following research objectives in the context of an online exam preparation product:

- Understand the efficacy of an AI tutor chatbot by residual score gain after controlling baseline ability.
- Identify the characteristics of student-AI interaction that are associated with residual gain.

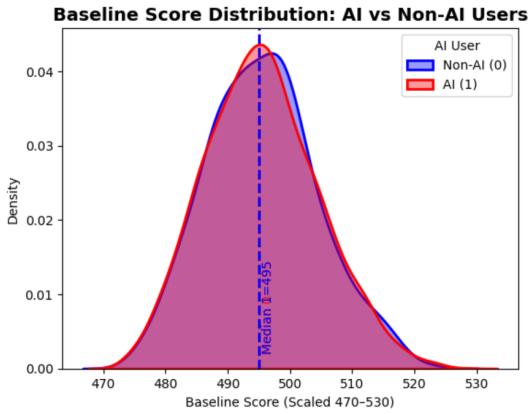


Figure 1: Pre-assessment score distributions for AI users ($n = 667$) and non-AI users ($n = 3{,}467$). The near-complete overlap indicates minimal selection bias on observed baseline ability.

- Identify learner engagement profiles which exhibit different residual gain.

2 Data

Chat history, course activity, and assessment data were collected from a cohort enrolled in an online pre-licensure exam preparation course between August 2024 and April 2025. The final sample comprises 667 AI tutor users (who interacted with the chatbot at least once) and 3,467 non-AI users (who did not interact with chatbot). Both groups took pre- and post-assessments to calculate scaled score gains. Pre-assessment score distributions were nearly identical across both groups (Figure 1), indicating minimal selection bias between the two groups.

3 Methods

The analysis proceeds through a four-step pipeline. Each step's output serves as input to the next; the sequential architecture is deliberate and has implications for model specification (see §3.1).

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 260–266
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

3.1 Residual Gain Modeling

Residual gain for AI user i is defined as:

$$r_i = g_i^{(\text{obs})} - \hat{g}(b_i; \hat{\theta}_{\text{non-AI}}) \quad (1)$$

where $g_i^{(\text{obs})}$ is the observed score gain, b_i is the baseline (pre-assessment) score, and $\hat{g}$ is a quadratic ordinary-least-squares model with parameters $\hat{\theta}_{\text{non-AI}}$ estimated exclusively from the non-AI cohort ($n = 3{,}467$). Functional form was selected by comparing linear and quadratic OLS specifications via a nested F-test. Distributional fairness of the residual was assessed by an ANOVA on residuals across baseline-score tertiles. Reference-model diagnostics are reported in §4.1.

The reference model deliberately conditions on baseline score. The other learning behavior variables (e.g., total questions answered, study intensity) are excluded because they serve as predictors in the downstream elastic-net decomposition (§3.3). Including them in both stages would constitute double adjustment: the residual is *designed* to contain all behavioral signal so that the elastic-net can allocate it. This follows standard practice in value-added modeling, where prior achievement is the primary conditioning variable and post-treatment covariates are excluded to avoid "bad controls" (Angrist and Pischke, 2009; Sanders and Horn, 1998; Raudenbush, 2004).

3.2 Cognitive Fingerprinting

To evaluate interaction quality, each student–AI conversation was classified using Gemini-1.5-Flash (temperature $= 0.1$) via a structured few-shot prompt, an established approach for text annotation (Gilardi et al., 2023; Ziems et al., 2024). Classifications spanned two dimensions:

- **Content section** (S): five exam topics (A–E).
- **Cognitive depth** (D): five levels adapted from Bloom's Taxonomy (Bloom, 1956)—Factual Recall, Conceptual Understanding, Application/Problem-Solving, Analysis/Reasoning, and Metacognitive/Strategic.

The Cartesian product $S \times D$ yields a 25-cell grid. Each student's proportion vector across this grid constitutes their **cognitive fingerprint**: a unique signature of how they engage with the AI across content areas and reasoning levels.

Classifier validation. Face validity was established through manual inspection of student prompt language: interactions classified as Analysis/Reasoning contained reasoning-oriented language (e.g., "why does X fail under Y"), while Factual Recall items contained definition-seeking queries (e.g., "what is X"). Convergent validity was assessed by examining within-student cross-section consistency: students who engaged at higher cognitive levels in one content section tended to do so across sections, a pattern inconsistent with random classification. Furthermore, a pilot inter-rater reliability design is planned in the future work (§6).

3.3 Feature Selection and Composite Indices

Twenty-five candidate predictors spanning four families—total course activity, cognitive depth (proportions across the five Bloom-adapted levels defined in §3.2), content topic, and student-AI interaction volume—were analyzed using elastic-net regression (Zou and Hastie, 2005) with residual gain as the outcome. Baseline score was included as a control. The elastic-net penalty combines an ℓ_1 component (which forces weak predictors to exactly zero, aligned with the feature-selection objective) and an ℓ_2 component (which stabilizes correlated features); the overall strength α and the ℓ_1/ℓ_2 mixing ratio were selected via 5-fold cross-validation, with fitted values reported in §4.2. The full predictor list, organized by family, appears in Appendix B.

Predictors with non-zero weights were retained and partitioned into two orthogonal composite indices: the **Effort Index** (course-activity features) and the **AI-Leverage Index** (AI-interaction features). These indices serve as inputs to the persona analysis in the clustering step.

3.4 Behavioral Clustering

A Gaussian Mixture Model (Reynolds, 2009) was fit on the Effort Index $\times$ AI-Leverage Index space. Baseline score was excluded so that personas reflect *behavior*, not ability. The optimal number of clusters was determined by comparing BIC, AIC, and Silhouette score across $k \in \{3, 4, 5\}$. Cluster stability was assessed via the Adjusted Rand Index (ARI) across 10 random initializations, and per-student assignment certainty was quantified using posterior class probabilities. Silhouette is interpreted as a secondary diagnostic given that it penalizes the soft, overlapping cluster boundaries

characteristic of GMMs and was originally developed for hard-partition methods such as k-means. Cluster diagnostics are reported in §4.3.

4 Results

4.1 Aggregate Effect

Reference-model diagnostics. The quadratic reference model ($R^2 = 0.053$) significantly outperformed a linear specification ($R^2 = 0.041$; nested $F = 44.26$, $p < 0.001$). The fitted equation is $\hat{g}(x) = -0.0066x^2 + 6.41x - 1544.38$; the negative leading coefficient captures the ceiling effect for high-baseline learners. An ANOVA on residuals across baseline tertiles was non-significant ($F = 2.98$, $p = 0.052$), indicating that residual gain does not systematically favor any ability stratum.

Aggregate residual gain. The mean residual gain across all 667 AI users is $+0.50$ points (directionally positive, not statistically significant). The median is exactly 0.00, with SD $= 6.22$ and a range from -10.54 to $+28.59$. This wide range and near-zero average suggest substantial heterogeneity and motivate the persona-level decomposition that follows. It indicate that *aggregate evaluation* is an insufficient paradigm when treatment effects are heterogeneous and likely depend on student-AI interaction quality.

4.2 Feature Selection

Model fit. Cross-validation selected $\alpha = 0.1851$ and ℓ_1-ratio $= 0.95$; the ℓ_2 component thus contributes only 5% of the total penalty, consistent with the feature-selection objective.

Selected predictors. Sixteen of 25 predictors maintained non-zero weights. Six features positively predicted residual gain (Figure 2): Total Questions Answered ($+1.65$), Study Efficiency ($+0.89$), Study Intensity ($+0.65$), Full-Length Test Effort ($+0.28$), Analysis/Reasoning% ($+0.23$), and Average Tokens per Interaction ($+0.02$).

The course-activity family yielded the largest weights, including Total Questions Answered, Study Efficiency, and Study Intensity. This confirms that increased practice alongside AI tutoring drives greater residual gains. Within the cognitive-depth dimension, notably, Analysis/Reasoning had a meaningful positive prediction, while Factual/Recall negatively predicted gain (-0.32).

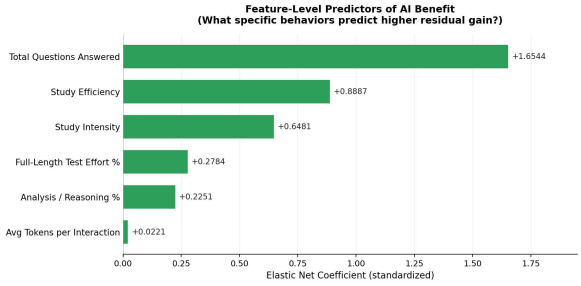


Figure 2: Standardized elastic-net coefficients for features positively predicting residual gain. All AI-volume metrics received zero weight (not shown).

Crucially, all AI interaction volume metrics received exactly zero weight. The number of chat sessions, student-AI interactions, and tokens contributed nothing to residual gain prediction once course effort and cognitive depth were in the model. This result is robust to reasonable perturbation of the elastic-net hyperparameters: volume features occupy a flat region of the coefficient surface with no predictive signal.

4.3 Learner Personas

Cluster diagnostics. All three model-selection criteria converged on $k = 4$: BIC dropped 540 points from $k = 3$ to $k = 4$ before rising at $k = 5$, AIC followed an analogous trajectory, and Silhouette peaked at $k = 4$ (Table 2 in Appendix). Under the final model, cluster assignments were highly stable across 10 random initializations (ARI $= 0.993 \pm 0.008$), and 96.7% of students were assigned to their persona with $> 80\%$ posterior probability (only 0.9% below 60%). The Silhouette score of 0.155 is low by conventional thresholds but expected given its sensitivity to the soft, overlapping cluster boundaries characteristic of GMMs; the personas nevertheless exhibit strong interpretive distinctness which is illustrated below.

Persona profiles. Four personas emerge from the GMM clustering (Table 1):

The cognitive fingerprints of these personas are qualitatively distinct (Figure 4). Maximizers exhibit a pronounced spike on Analysis/Reasoning; Grinders collapse almost entirely into Factual Recall despite comparable course effort; Satisficers show a balanced mid-depth profile; Sideliners are undifferentiated across all dimensions. This interpretive distinctness complements the statistical stability reported above and addresses concerns about geometric cluster separation.

Persona	n	%	Cognitive Signature	Residual Gain
Maximizer	25	3.7	Analysis + Application	$+2.82$
Satisficer	255	38.2	Conceptual + Application	$+1.00$
Grinder	43	6.4	Factual Recall	$+0.77$
Sideliner	344	51.6	Shallow, undifferentiated	-0.07

Table 1: Behavioral definition, cognitive signature, and mean residual gain of the four learner personas.

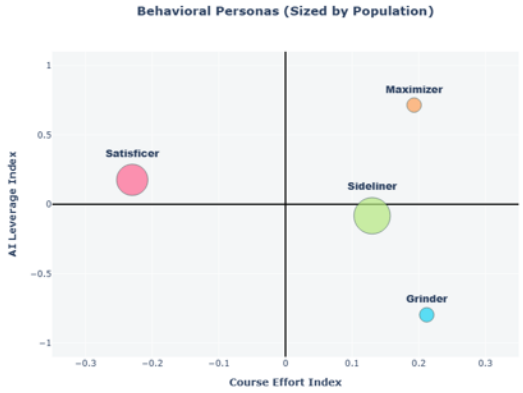


Figure 3: Learning personas in the Course Effort $\times$ AI Leverage plane (bubble size proportional to n-count).

The Maximizer–Grinder comparison. Figure 3 present the four learning personas across the two dimensions: Course Effect (x-axis) and AI-Levarage (y-axis).It is worth noting that both Maximizers and Grinders share high course effort but differ sharply in AI-Leverage. The residual gain gap of $+2.05$ points ($+2.82$ vs. $+0.77$) isolates the association between AI interaction cognitive depth and learning gain, net of effort. This within-treatment comparison holds the primary confounder (course effort) approximately constant while varying the dimension of interest (cognitive depth), providing the strongest quasi-experimental evidence available in this observational design.

Figure 4 further highlights the distinct cognitive-depth fingerprints of the four personas. Maximizers were uniquely characterized by substantial engagement in Analysis/Reasoning. Conversely, Grinders collapsed almost entirely into Factual/Recall despite high course effort, while Satisficers and Sideliners clustered centrally with low-depth engagement.

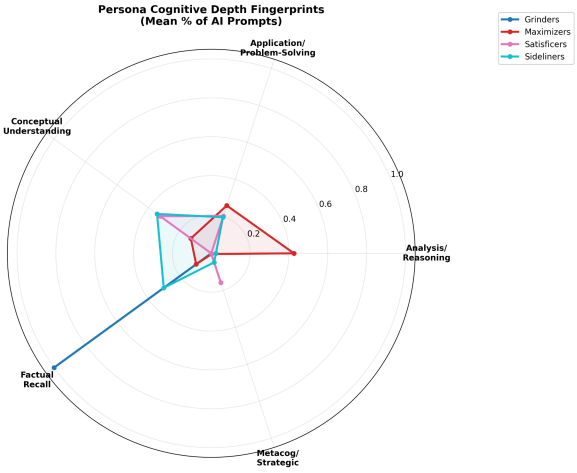


Figure 4: Mean cognitive depth distribution of AI interactions by persona. Maximizers spike on Analysis/Reasoning; Grinders collapse to Factual Recall.

5 Conclusion

This study demonstrates that the efficacy of AI tutoring is driven by the cognitive depth of student-AI interaction rather than sheer usage volume. Learners engaging in higher-order thinking (especially Analysis and Reasoning) achieved greater residual gains, whereas those focused primarily on Factual Recall derive negligible benefit.

The finding aligns with Bastani et al. (2025), who showed that unguided ChatGPT improved practice performance but harmed actual learning, while a purpose-built tutor with guardrails, especially providing scaffolding hints instead of giving away answers, avoided this harm. Our Maximizer profile represents the behavioral signature of deep, guided-like AI engagement; the Sideliner profile represents unguided, shallow use.

The non-significant aggregate effect ($+0.50$) is largely explained by population composition: 51.6% of AI users are Sideliners (gain $\approx$ 0), while only 3.7% are Maximizers (gain $= +2.82$). Aggregate efficacy metrics will systematically miss persona-level variation, and product decisions based on aggregate effects alone risk missing the important engagement factors driving learning gains.

By shifting the evaluation of AI tools from quantitative usage to qualitative cognitive fingerprints, this research provides vital evidence that student-AI interfaces should be optimized with Socratic scaffolding and deep-thinking prompts to foster deep cognitive engagement. Further research should focus on validating these findings across

broader academic domains and diverse student populations.

6 Future Work

To further strengthen the study, the following extensions are planned as immediate next step to validate the LLM-as-a-judge classification framework of cognitive depth.

- **Inter-rater reliability** A stratified sample of 50 interactions (8 per cognitive depth level, stratified by Gemini's assigned label, plus 10 randomly drawn from the full pool) will be independently coded by human subject matter experts blind to the model's classifications.The human experts will use the identical rubric definitions and few-shot examples supplied to the Gemini classifier. Cohen's κ, per-class agreement, and a confusion matrix will be computed against the human experts. Both inter-rater agreement and human-model agreement will be evaluated.

Limitations

This study is quasi-experimental: students were not randomly assigned to AI usage or to specific engagement patterns, so unobserved confounders (e.g., metacognitive awareness, intrinsic motivation) cannot be ruled out despite baseline-ability adjustment and the effort-matched Maximizer–Grinder contrast. Furthermore, statistical power for the Maximizer persona is also limited ($n = 25$); the $+2.82$ residual-gain estimate is stable across cluster reassignments (96.7% posterior certainty) but derives from a small subgroup.

A future randomized trial study assigning students to scaffolded versus unscaffolded and unguided AI interaction, and replicated in a larger cohort, would strengthen the conclusion from current quasi-experimental study to a causal estimate of depth-of-engagement effects.

Finally, findings were derived from a single exam preparation product; cross-domain validation is needed to assess whether the same learning persona structure and AI interaction depth-over-volume result will generalize in other domains and subjects.

References

Joshua D. Angrist and Jörn-Steffen Pischke. 2009. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press.

Hamsa Bastani, Osbert Bastani, Alp Sungu, Haosen Ge, Özge Kabakcı, and Rei Mariman. 2025. Generative AI without guardrails can harm learning: Evidence from high school mathematics. *Proceedings of the National Academy of Sciences*, 122(26):e2422633122.

Benjamin S. Bloom. 1956. *Taxonomy of Educational Objectives: The Classification of Educational Goals*. David McKay Company.

Fabrizio Gilardi, Meysam Alizadeh, and Maël Kubli. 2023. ChatGPT outperforms crowd workers for text-annotation tasks. *Proceedings of the National Academy of Sciences*, 120(30):e2305016120.

Greg Kestin, Kelly Miller, Anna Klales, Timothy Milbourne, and Giada Ponti. 2024. AI tutoring outperforms active learning. *Harvard Initiative for Learning and Teaching Working Paper*.

Ethan Mollick and Lilach Mollick. 2024. Instructors as innovators: A future-focused approach to new AI learning opportunities, with prompts. *Wharton School Working Paper*.

Radek Pelánek, Jiří Řihák, and Jan Papoušek. 2016. Impact of data collection on interpretation and evaluation of student models. In *Proceedings of the Sixth International Conference on Learning Analytics & Knowledge*, pages 40–49. Association for Computing Machinery.

Stephen W. Raudenbush. 2004. What are value-added models estimating and what does this imply for statistical practice? *Journal of Educational and Behavioral Statistics*, 29(1):121–129.

Douglas A. Reynolds. 2009. Gaussian mixture models. In Stan Z. Li and Anil Jain, editors, *Encyclopedia of Biometrics*, pages 659–663. Springer.

William L. Sanders and Sandra P. Horn. 1998. Research findings from the Tennessee Value-Added Assessment System (TVAAS) database: Implications for educational evaluation and research. *Journal of Personnel Evaluation in Education*, 12(3):247–256.

William R. Shadish, Thomas D. Cook, and Donald T. Campbell. 2002. *Experimental and Quasi-Experimental Designs for Generalized Causal Inference*. Houghton Mifflin.

Caleb Ziems, William Held, Omar Shaikh, Jiaao Chen, Zhehao Zhang, and Diyi Yang. 2024. Can large language models transform computational social science? *Computational Linguistics*, 50(1):237–291.

Hui Zou and Trevor Hastie. 2005. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2):301–320.

k	BIC (↓)	AIC (↓)	Silhouette (↑)
3	1042.1	965.5	0.097
4	**502.0**	**398.5**	**0.155**
5	537.7	407.1	0.099

Table 2: Model-selection criteria across candidate cluster counts. All three criteria identify $k = 4$ as optimal.

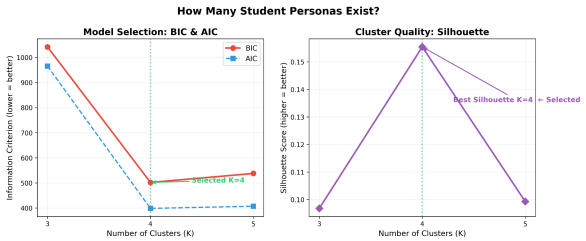


Figure 5: Cluster model selection. BIC and AIC (left) both minimize at $k = 4$; Silhouette (right) peaks at $k = 4$. All three criteria converge on four behavioral personas.

A Model Selection for Clustering

In Table 2, the BIC dropped 540 points from $k = 3$ to $k = 4$ before rising at $k = 5$; AIC followed an analogous trajectory. In Figure 5 Silhouette peaked at $k = 4$. The convergence of three independent criteria substantially reduces the risk that the four-cluster solution is an artifact of any single metric.

B Candidate Predictors

Table 3 lists all candidate predictors submitted to the elastic-net model, organized by family.

Family	Predictor	Definition
Course Activity	Total Questions Answered	Raw count of all practice questions answered
	Study Intensity	Total study hours / (study-window days + 1)
	Study Efficiency	Total questions answered ÷ total study hours
	Full-Length Test Effort %	Share of total questions from full-length practice tests
	Quiz Effort %	Proportion of effort directed toward quizzes
	Question Effort %	Proportion of effort directed toward question solving
	Full-length Attempt %	Share of practice sessions that were full-length tests
	Quiz Attempt %	Proportion of attempts focused on Quiz practice
	QBank Attempt %	Proportion of attempts focused on QBank practice
	Questions per Quiz	Average questions per Quiz session
	Questions per QBank	Average questions per QBank session
Cognitive Depth	Analysis/Reasoning %	Proportion of AI interactions classified as Analysis/Reasoning
	Application %	Proportion classified as Application/Problem-Solving
	Conceptual Understanding %	Proportion classified as Conceptual Understanding
	Factual Recall %	Proportion classified as Factual Recall
	Metacognitive/Strategic %	Proportion classified as Metacognitive/Strategic
	Human-to-AI Token Ratio	Mean ratio of student to AI tokens per interaction
	Contextual / Non-Content %	Proportion of AI interactions classified as Contextual/Non-Content (non-academic or logistical)
Topic	Section Focus % (A–E)	Proportion of AI interactions per content from A to E
AI Volume	AI Usage Volume	Standardized composite of session/interaction/token counts
	Avg Tokens per Interaction	Mean total tokens per interaction

Table 3: Full list of candidate predictors for elastic-net feature selection.

Automatic Prompt Engineering for Generative AI–Based Essay Scoring

Yue Huang
Measurement Incorporated
yueh@udel.edu

Abstract

This study examined *automatic prompt engineering* (APE) for holistic scoring of one independent argumentative writing item from the PERSUADE 2.0 dataset. The APE approach optimized scoring instructions and examples for a GPT-4–family model under different configurations. The finalized configuration, selected through validation accuracy, was compared with a research-informed, zero-shot Baseline prompting approach on the held-out test set. Quadratic-weighted kappa (QWK) increased from .646 to .812, alongside improved exact agreement, adjacent agreement, and mean-score alignment. Gains were concentrated at higher score levels. Gender gaps narrowed descriptively, while English language learner subgroup patterns varied across metrics. These findings support prompt optimization for *automated essay scoring* (AES) while highlighting the importance of rubric review and subgroup evaluation.

1 Introduction

Automated essay scoring (AES) uses computational methods to assign essay scores that approximate those awarded by human raters. Its development has progressed from models based on engineered linguistic features to neural approaches that learn representations of writing, and more recently to generative language models (Huang et al., 2025). Rapid advances in generative artificial intelligence (AI) have renewed interest in using these models for AES. Their ability to interpret natural-language instructions offers a way to specify scoring tasks without training a separate scoring model for each application.

Recent research suggests that generative AI can approximate human essay ratings through prompt engineering, the deliberate design of model inputs to guide the content and format of its scores (Huang and Wilson, 2025; Huang et al., 2025, 2026). In AES, these inputs may specify a rater role, the writing task, scoring criteria, and examples of human-scored essays. Much of the prompt-based research reviewed by Huang et al. (2025) relied on researcher-designed prompts. Their synthesis identified useful design strategies, including clear instructions and rubric guidance, while also showing that the benefits of individual components varied across studies. Such guidance supports prompt construction, but leaves researchers to determine how components should be worded and combined for a particular scoring task.

Automatic prompt engineering (APE), a form of prompt optimization, uses algorithms to generate and evaluate alternative instructions and examples. Huang et al. (2025) recommended investigating DSPy, a framework for specifying language-model tasks and optimizing their prompts, for AES. Here, prompt optimization refers to searching for a prompt that improves agreement with human ratings; it does not update the language model's parameters. Examining both prompt content and scoring performance is important for understanding what this process contributes to AES.

2 Related Work

Prior research identifies two connected questions in prompt development: which information to provide and how to select an effective formulation. The AES studies synthesized by Huang et al. (2025) addressed the first through variations in rubric guidance, context, and examples. Huang et al. (2026) further showed that contextual information and scored examples influenced GPT-4o scoring of elementary students' essays. Together, these studies provide evidence about which information should be included in scoring prompts. However, they offer less guidance on how prompt formulations can be systematically generated and refined for a new scoring task beyond comparing a small set of

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 267–274
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

manually specified alternatives.

Research on prompt optimization addresses this problem by treating prompts as candidates whose performance can guide further search. Zhou et al. (2023) generated and evaluated multiple instruction candidates, with an iterative variant further exploring alternative formulations of high-performing candidates. Later approaches made revision more explicitly feedback-driven: ProTeGi (Prompt Optimization with Textual Gradients) used natural-language analyses of model errors to identify weaknesses in the current prompt and guide its iterative refinement, whereas OPRO (Optimization by PROmpting) used histories of candidate prompts and their scores (Pryzant et al., 2023; Yang et al., 2024). These methods differ in how they propose revisions, but share the principle that prompt development can be guided by observed performance rather than researcher judgment alone.

This principle extends beyond instruction wording to demonstrations, examples included in a prompt to illustrate the task. Wan et al. (2024) showed that exemplar selection can be as consequential as instruction optimization in the tasks they examined, motivating attention to what a model is shown as well as what it is told. Extending automatic prompt optimization to multi-stage language-model programs, Opsahl-Ong et al. (2024) brought these elements together in MIPRO (Multi-prompt Instruction PRoposal Optimizer), which jointly optimizes instructions and demonstrations across the connected LLM calls. For AES, this creates an opportunity to adapt both the description of the task and the examples illustrating its application, although the effectiveness of the resulting combination still requires evaluation in the target assessment context.

Educational applications of APE provide promising but distinct precedents. Xue et al. (2025) applied APE to 11 science constructed-response items, evaluating optimized scoring prompts in zero-shot and few-shot settings and examining how prompt changes affected scoring accuracy and agreement with human scores. Harada et al. (2026) extended prompt optimization to essay-scoring rubrics, iteratively refining them based on discrepancies between model-generated and human scores and the model's scoring rationales. These studies establish that prompt optimization is relevant to educational scoring, while addressing different parts of the scoring prompt. Xue et al. (2025) retained fixed examples within their few-shot conditions, and Harada et al. (2026) focused on rubric revision. They therefore leave room to examine joint instruction and demonstration optimization for holistic essay scoring.

The writing task also gives this extension substantive importance. Xue et al.'s (2025) short science responses were evaluated against item-specific content criteria. In the independent argumentative essays examined in this study, students develop a position through their own reasons and examples, and holistic scoring integrates the quality of development, organization, and language across an extended text. Effective responses can follow different argumentative paths rather than reproduce a single expected answer. This distinction does not imply that short-response scoring is simple; it means that findings from content-focused responses cannot establish how optimized prompts will interpret broader writing-quality criteria. Harada et al.'s (2026) essay study makes this a question about the optimization approach within AES, rather than an entirely unexamined application.

Evaluation also needs to consider whom a scoring approach serves. Huang et al. (2026) found that overall scoring performance and subgroup patterns did not necessarily move together. A brief comparison by gender and English language learner (ELL) status in this present study can therefore complement aggregate results by identifying whether agreement patterns differ across groups. Such comparisons provide preliminary evidence about the consistency of scoring performance, rather than a comprehensive test of fairness.

3 The Present Study

This study examines whether an APE approach improves the generative AI–based holistic scoring of independent argumentative essays relative to a researcher-developed zero-shot prompt informed by prior research on effective prompt-based AES strategies, hereafter termed the Baseline prompting approach. It extends prior educational prompt optimization research by jointly optimizing instructions and demonstrations (examples) and evaluating prompt content and agreement with human ratings. An exploratory analysis also examines performance across gender and ELL groups. The findings can inform scalable, rubric-aligned scoring systems by evaluating a systematic alternative to manual prompt development. Three research questions guide the study:

RQ1. How does the prompt generated by the APE approach differ in content from that used by the Baseline prompting approach?

RQ2. How do the APE approach and Baseline prompting approach compare in agreement with human ratings and score-level prediction patterns?

RQ3. How does scoring performance for the two approaches vary across gender and ELL groups?

4 Methods

4.1 Sample

This study used PERSUADE 2.0, a publicly available dataset of argumentative essays written by US students in Grades 6–12 (Crossley et al., 2024). The analysis was restricted to one independent-writing assignment asking students to take a position on an argumentative topic and support it with reasons and examples. Responses were retained when essay and assignment text were nonempty and the human holistic score was valid. The resulting 2,148 essays came from Grades 9–12, with 12, 191, 1,548, and 397 responses, respectively. Female students contributed 48.8% of essays, and ELL students contributed 27.9%. Dataset-reported essay lengths ranged from 159 to 1,584 words. Human raters scored the essays using the six-point holistic rubric for independent writing provided with the dataset, with 1 indicating the lowest performance level and 6 the highest.

Following the competition split recorded in the dataset, I allocated 1,492 essays (69.5%) to a training set for prompt development and 656 essays (30.5%) to a held-out test set for evaluating scoring performance. The training set was further divided into separate subsets for developing candidate prompts and validating their performance, as detailed in *Section 4.3*. The test essays were excluded from both prompt optimization and configuration selection and used only to evaluate and compare the APE approach and Baseline prompting approach.

4.2 Prompting Approaches

Both approaches scored essays through the GPT application programming interface (API). Scoring refers to applying a specified prompt to an essay and requesting one holistic score within the rubric's six-point scale. The APE approach used the MIPROv2 optimizer within DSPy framework to jointly optimize scoring instructions and essay–score demonstrations for agreement with human scores (Opsahl-Ong et al., 2024; DSPy, 2026).

The Baseline prompting approach used a researcher-written zero-shot prompt. Drawing on Huang and Wilson (2025) and Huang et al. (2026), it included an expert-rater role, the scoring task, assignment and student context, the original six-level holistic rubric, and a requirement to return one integer from 1 to 6 to reflect the writing quality following the rubric. This instruction also served as the starting point for the APE approach. During optimization, DSPy could rewrite the instruction and select essay–score demonstrations. The comparison therefore evaluated the complete APE approach against the Baseline prompting approach, rather than isolating individual prompt components.

4.3 Prompt Optimization and Configuration

To develop the APE approach, I implemented MIPROv2 in DSPy version 3.2.1 (DSPy, 2026). MIPROv2 is an updated implementation of the MIPRO method described by Opsahl-Ong et al. (2024), retaining its core strategy of jointly optimizing instructions and demonstrations through Bayesian search. Within each configuration, training essays were divided into candidate-prompt-development and validation subsets. Using the Baseline instruction and candidate-prompt-development data, the optimizer proposed alternative instructions and demonstration sets. Each search trial evaluated a candidate prompt by scoring the full validation subset and calculating exact agreement with human ratings. Completed trials guided subsequent search, and the candidate prompt with the highest validation agreement was retained.

Validation-set size and search budget varied across optimization configurations. Larger validation subsets provided more evidence for evaluating prompts, while additional trials explored more prompt combinations; both increased the number of scoring requests. These settings explored practical development choices rather than demonstrating cost savings or establishing a universally optimal configuration.

Table 1 summarizes the configurations. Configuration 2 increased validation-set size relative to Configuration 1, whereas Configuration 3 increased the search budget. Configurations 1 and 3 shared the same 100 validation essays, included within Configuration 2's 200-essay subset. Each configuration was optimized separately. GPT-4.1-mini proposed prompting instructions at an initial tem-

perature of 1.0. Configuration 3 had the highest validation exact agreement (.540) and was selected for the APE approach reporting. Its optimized instruction and six demonstrations were finalized before held-out evaluation.

Config.	Candidate Development (*N*)	Validation (*N*)	Search Trials (*N*)	Validation Exact Agreement
1	1,392	100	12	.510
2	1,292	200	12	.525
3	1,392	100	20	.540

Table 1: Configurations for the APE approach.

4.4 Essay Scoring Procedure

After finalizing the prompting approaches, I applied each to all held-out essays using GPT-4.1-mini through the API, with scoring temperature set to 0 to reduce sampling variability. Each request supplied one target essay and requested a single holistic score. The Baseline prompting approach used the zero-shot researcher-written instruction; the APE approach used the finalized optimized instruction and demonstrations.

4.5 RQ1 and RQ2: Prompt Content and Scoring Performance Evaluation

For RQ1, I compared the final prompts in instruction content, rubric presentation, and demonstrations. For RQ2, I evaluated human–machine agreement using quadratic-weighted kappa (QWK), exact agreement, adjacent agreement, the proportion within one point; and standardized mean difference (SMD). QWK ≥ 0.70 and $|\text{SMD}| \leq 0.15$ served as reference criteria from Williamson et al. (2012).

Both approaches used the same valid human–machine score pairs. Differences were calculated as APE minus Baseline. I also estimated 95% percentile confidence intervals from 10,000 paired bootstrap resamples.

Finally, confusion matrices displayed counts and percentages within each human score level to examine where the approaches agreed with human ratings and where they under- or over-scored.

4.6 RQ3: Subgroup Comparisons

For RQ3, I calculated the four metrics separately by gender and ELL status and descriptively compared subgroup performance under the Baseline prompting approach and the APE approach. Gender and ELL contrasts were calculated as male minus female and non-ELL minus ELL, respectively, for QWK, exact agreement, and adjacent agreement. For SMD, I reported absolute values within groups and absolute differences between groups' absolute SMD values.

5 Results

5.1 RQ1: Comparison of Prompt Content

The APE approach's final instruction retained the Baseline prompting approach's expert-rater role, task context, six-point scale, and numeric-score requirement. It condensed the instruction from over 700 to just over 100 words, excluding demonstrations and the target essay. The Baseline prompt reproduced the assignment and all six rubric descriptors. The exact original rubric was not presented in the APE prompt. Instead, the APE prompt summarized dimensions from the rubric including argument development, critical thinking, organization, language, and mechanics.

The APE approach added six examples in its demonstrations with scores of 5, 3, 5, 5, 4, and 5, including one bootstrapped and five labeled examples. These did not cover the full scale. The shorter instruction did not imply a shorter complete prompt because example essays added text. The comparison therefore combined changes in wording, rubric presentation, and demonstrations.

5.2 RQ2: Comparison of Scoring Performance

The APE approach showed higher overall agreement with human ratings than the baseline prompting approach (Table 2). QWK increased from .646, below the .70 reference threshold, to .812, above it—a difference of .166. Exact agreement increased from 39.9% to 54.6%, a gain of 14.6 percentage points; adjacent agreement increased from 90.4% to 98.8%, a gain of 8.4 points. The paired intervals for all three differences excluded zero.

The human mean was 4.209, compared with 3.651 for the Baseline prompting approach and 4.331 for the APE approach. The Baseline prompting approach under-scored on average (SMD = $-.494$), whereas the APE approach slightly over-scored (.108). Absolute SMD therefore decreased. The APE approach's absolute SMD met the .15 criteria. The signed difference (.602) reflects the shift from under-scoring to over-scoring.

Metric	Baseline	APE	Difference
QWK	0.646	0.812	$+0.166$ $[0.137, 0.196]$
Exact agreement	39.9%	54.6%	$+14.6$ pp $[9.0, 20.4]$
Adjacent agreement	90.4%	98.8%	$+8.4$ pp $[6.1, 10.7]$
SMD	-0.494	0.108	$+0.602$ $[0.556, 0.651]$

Table 2: Human–machine agreement and paired differences on the test set.
Note. $N = 656$. Baseline = the Baseline prompting approach; APE = the APE approach. Difference = APE − Baseline. Brackets contain 95% confidence intervals. Agreement differences are percentage points (pp). SMD values and their difference are signed.

5.3 RQ2: Comparison of Predictions Across Score Levels

Figure 1 shows that improvement varied across score levels. Of 86 human-rated score-6 essays, the Baseline prompting approach assigned 6 to one (1%), whereas the APE approach assigned 6 to 54 (63%). The APE approach substantially reduced upper-end under-scoring, while also assigning 6 to 63 essays rated 5 and three rated 4 by human raters. Consequently, 54 of its 120 score-6 predictions (45%) matched the human rating.

At human score 1, the Baseline prompting approach matched two of three essays, whereas the APE approach assigned 2 to all three. Thus, the APE predictions remained within one point, although exact agreement decreased. Only three essays supported this comparison, limiting conclusions at the lowest score level; the score-6 pattern had a stronger descriptive basis (N = 86).

For human scores 2–5, exact agreement was 55%, 53%, 42%, and 42% for the Baseline prompting approach, compared with 55%, 50%, 62%, and 47% for the APE approach. Gains were concentrated at scores 4–6. Overall, under-scored essays decreased from 347 to 110, while over-scored essays increased from 47 to 188. The APE approach reduced the Baseline's downward scoring tendency without improving every score level.

5.4 RQ3: Comparison Across Student Groups

The APE approach showed higher agreement and lower absolute SMD in every subgroup (Table 3). Gender gaps were descriptively smaller with the APE approach than with the Baseline prompting approach across all four metrics. The APE approach achieved QWK above .80 for both genders, with absolute SMD of .102 for male and .116 for female students. These point estimates met both reference thresholds.

Across ELL groups, the QWK gap decreased from .118 under the Baseline prompting approach to .029 under the APE approach, although neither group's QWK reached .70 under the APE approach. Under the Baseline approach, exact agreement was higher for non-ELL students; under the APE approach, it was higher for ELL students. With the APE approach, adjacent agreement was nearly 99% for both groups. Although absolute SMD decreased for both groups under the APE approach, the difference between the groups increased from .097 under the Baseline approach to .176 under the APE approach. The APE approach met the .15 criterion for ELL students (.016), while some over-scoring remained for non-ELL students (.192). These comparisons describe observed subgroup patterns rather than statistically established differences.

6 Discussion

This study demonstrates the promise of automatic prompt engineering (APE) for improving generative AI–based essay scoring. By connecting prompt-content analysis with held-out performance and subgroup comparisons, it shows how prompt optimization can strengthen agreement with human ratings while identifying priorities for further refinement. The findings support a systematic approach to prompt development that combines research-informed design with empirical evaluation.

An important finding is that the APE approach produced a condensed instruction paired with scored demonstrations, suggesting that scoring expectations can be communicated through complementary forms of guidance: instructions identify relevant qualities, while examples illustrate their application. This connects with Xue et al.'s (2025) observation that APE can reformulate scoring instructions, raising questions about whether optimization changes the intended scoring criteria. In the present study, stronger human–machine agreement suggests that concise instructions combined with examples can effectively support scoring, although agreement alone does not establish that all rubric criteria were preserved. The key distinction is therefore between refining how criteria are communicated and changing what is assessed. Expert

review against the original rubric could help maintain this distinction, while comparisons of full and summarized rubrics, with and without demonstrations, could clarify the contributions of instruction changes and examples.

Another key finding is the APE approach's substantial improvement in scoring performance over an already well-structured baseline prompting approach. Its rater role, assignment context, complete rubric, and output requirements reflected prior research guidance. The gains therefore demonstrate the added value of automatic optimization beyond a carefully designed zero-shot starting point. Huang et al. (2026) highlighted the benefits of contextual information and scored examples for prompt-based essay scoring. Building on this evidence, the present study shows how APE can use human-scored data to refine scoring instructions and select essay–score demonstrations, translating general prompting guidance into a configuration tailored to the scoring task. This finding supports treating research-informed prompt design as a foundation for further empirical optimization and positions APE as a systematic means of improving agreement with human scores.

A particularly relevant comparison comes from Plasencia-Calaña (2025), who evaluated the same assignment within PERSUADE 2.0. Reported QWK ranged from .48 for zero-shot prompting and .60 for chain-of-thought with few-shot examples to .77 for the best-performing feature-based histogram gradient boosting classifier. The present APE approach's estimate of .812 is numerically higher, providing encouraging context for prompt optimization without model fine-tuning. However, differences in models and evaluation procedures preclude attributing this cross-study difference solely to APE. Within the present comparison, improved recognition of highly rated essays further demonstrates its practical value, while additional highest-score assignments identify opportunities for calibration.

Subgroup findings also highlight the importance of examining how gains are distributed. The APE approach produced similar performance across gender groups, whereas ELL patterns differed by metric. Wilson and Huang (2024) found that both human and automated scores were less predictive of state writing-test performance for ELL students, situating these concerns within broader writing assessment. In the present study, differences in human-score distributions between ELL and non-ELL students may also contribute to the observed metric differences, particularly because the APE approach's performance varied across score levels.

Overall, this pilot study extends educational prompt optimization research by examining the APE approach for scoring full-length independent argumentative essays. The findings show that jointly optimizing instructions and demonstrations can strengthen research-informed scoring prompts and improve agreement with human ratings. Integrating such data-guided refinement with rubric review and subgroup evaluation provides a constructive foundation for developing scalable AES systems and strengthening the evidence supporting their use.

6.1 Limitations and Future Research

First, the study was limited to one assignment, one generative AI model, and one researcher-written baseline prompt. Future work should examine generalizability across writing tasks, models, and baseline prompts, and assess optimization stability using different seeds and validation samples.

Second, configuration comparisons used partly overlapping validation subsets of different sizes and did not establish an optimal search budget. More controlled comparisons, alongside direct measurements of token use and processing time, would clarify the balance between scoring performance and computational demands.

Third, sparse lowest-score observations and the descriptive nature of subgroup analyses limit conclusions about score-level and demographic differences. Larger samples at underrepresented score levels and confirmatory subgroup analyses would strengthen these comparisons.

Lastly, the Baseline prompting approach was research-informed but zero-shot, so gains from the APE approach cannot be attributed separately to instruction optimization and added demonstrations. Comparisons with Plasencia-Calaña (2025) provide encouraging context, but differing evaluation procedures limit direct comparisons. Comparing the APE approach with a research-informed few-shot approach would clarify its added value.

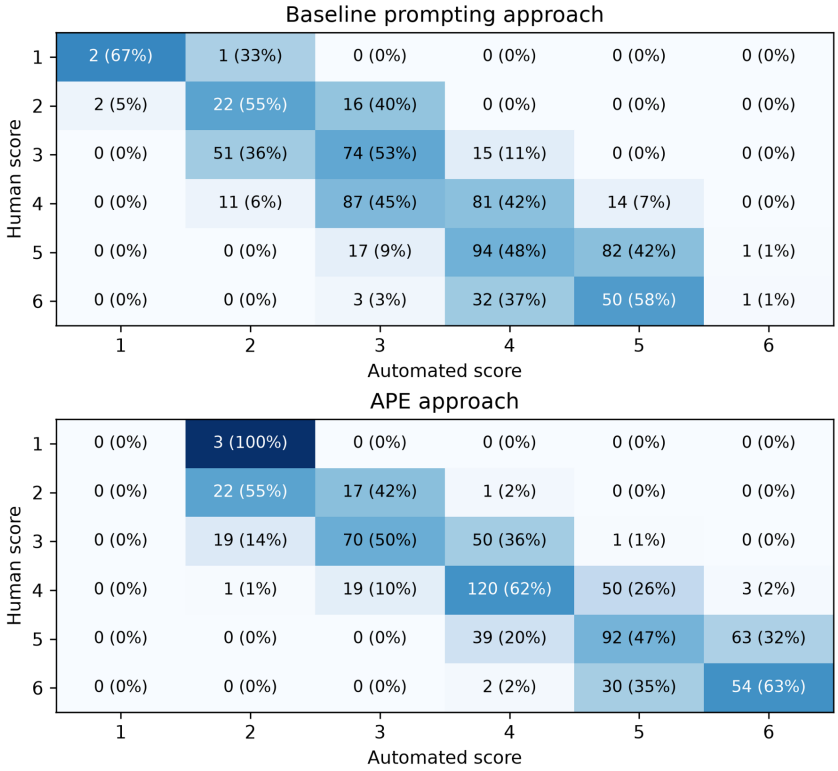


Figure 1: Confusion matrices for the Baseline prompting approach and the APE approach.

Note. Rows indicate human scores and columns automated scores. Cells show essay counts and percentages within each human-score row. Human-score counts for levels 1–6 were 3, 40, 140, 193, 194, and 86, respectively.

Group/Contrast	*N*	Approach	QWK	Exact Agreement	Adjacent Agreement	Absolute SMD*
Male	358	Baseline	.659	41.6%	91.9%	.458
Female	298	Baseline	.630	37.9%	88.6%	.538
Male − Female	—	Baseline Δ	.029	3.7	3.3	.080
Male	358	APE	.815	55.6%	98.9%	.102
Female	298	APE	.806	53.4%	98.7%	.116
Male − Female	—	APE Δ	.009	2.2	0.2	.014
Non-ELL	429	Baseline	.475	41.7%	89.5%	.590
ELL	227	Baseline	.357	36.6%	92.1%	.687
Non-ELL − ELL	—	Baseline Δ	.118	5.2	−2.6	.097
Non-ELL	429	APE	.695	52.2%	98.6%	.192
ELL	227	APE	.666	59.0%	99.1%	.016
Non-ELL − ELL	—	APE Δ	.029	−6.8	−0.5	.176

Table 3: Scoring performance by gender and ELL status.

Note. Baseline = the Baseline prompting approach; APE = the APE approach. Exact and adjacent agreement differences are expressed in percentage points. *Group rows report absolute SMD; contrast rows report the magnitude of the between-group difference in absolute SMD.

References

Scott Andrew Crossley, Yu Tian, Perpetual Baffour, Alex Franklin, Meg Benner, and Ulrich Boser. 2024. A large-scale corpus for assessing written argumentation: PERSUADE 2.0. *Assessing Writing*, 61:100865.

DSPy. 2026. DSPy source code (version 3.2.1). GitHub. Version 3.2.1, released May 5, 2026.

Keno Harada, Lui Yoshida, Takeshi Kojima, Yusuke Iwasawa, and Yutaka Matsuo. 2026. Automated refinement of essay scoring rubrics for language models via reflect-and-revise. In *Proceedings of the 30th Conference on Computational Natural Language Learning*, pages 771–789, San Diego, California, USA. Association for Computational Linguistics.

Yue Huang, Corey Palermo, Ruitao Liu, and Yong He. 2025. An early review of generative language mod-

els in automated writing evaluation: Advancements, challenges, and future directions for automated essay scoring and feedback generation. *Chinese/English Journal of Educational Measurement and Evaluation*, 6(2):5.

Yue Huang, Corey Palermo, and Joshua Wilson. 2026. Accuracy and fairness of generative AI in automated essay scoring: Comparing GPT-4o, feature-based models, and human raters. *Assessing Writing*, 69:101047.

Yue Huang and Joshua Wilson. 2025. Evaluating LLM-based automated essay scoring: Accuracy, fairness, and validity. In *Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con): Works in Progress*, pages 71–83, Pittsburgh, Pennsylvania, United States. National Council on Measurement in Education (NCME).

Krista Opsahl-Ong, Michael J. Ryan, Josh Purtell, David Broman, Christopher Potts, Matei Zaharia, and Omar Khattab. 2024. Optimizing instructions and demonstrations for multi-stage language model programs. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 9340–9366, Miami, Florida, USA. Association for Computational Linguistics.

Yenisel Plasencia-Calaña. 2025. Operationalizing automated essay scoring: A human-aware approach. *arXiv*, 2506.21603. Preprint.

Reid Pryzant, Dan Iter, Jerry Li, Yin Tat Lee, Chenguang Zhu, and Michael Zeng. 2023. Automatic prompt optimization with "gradient descent" and beam search. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 7957–7968, Singapore. Association for Computational Linguistics.

Xingchen Wan, Ruoxi Sun, Hootan Nakhost, and Sercan Ö. Arık. 2024. Teach better or show smarter? on instructions and exemplars in automatic prompt optimization. In *Advances in Neural Information Processing Systems*, volume 37, pages 58174–58244. Neural Information Processing Systems Foundation, Inc. (NeurIPS).

David M. Williamson, Xiaoming Xi, and F. Jay Breyer. 2012. A framework for evaluation and use of automated scoring. *Educational Measurement: Issues and Practice*, 31(1):2–13.

Joshua Wilson and Yue Huang. 2024. Validity of automated essay scores for elementary-age English language learners: Evidence of bias? *Assessing Writing*, 60:100815.

Mingfeng Xue, Yunting Liu, Xingyao Xiao, and Mark Wilson. 2025. Automatic prompt engineering for automatic scoring. *Journal of Educational Measurement*, 62(4):559–587.

Chengrun Yang, Xuezhi Wang, Yifeng Lu, Hanxiao Liu, Quoc V. Le, Denny Zhou, and Xinyun Chen. 2024. Large language models as optimizers. In *The Twelfth International Conference on Learning Representations*.

Yongchao Zhou, Andrei Ioan Muresanu, Ziwen Han, Keiran Paster, Silviu Pitis, Harris Chan, and Jimmy Ba. 2023. Large language models are human-level prompt engineers. In *The Eleventh International Conference on Learning Representations*.

A Comprehensive Evaluation of GenAI-based Items in a Medical Examination

Yanlin Jiang, Marcus Walker, Andrew Dallas,
Aquia Richburg, Nikole Gregg, Brittany Corrigan

National Commission on Certification of Physician Assistants

Abstract

This study evaluated GenAI-based medical assessment items against subject matter expert developed items. Analyses of item content and response data from a high-stakes examination showed comparable content quality and psychometric characteristics across the two item types. These findings provide empirical support for the quality of GenAI-based items and their potential use in educational and professional assessment programs.

1 Introduction

As generative artificial intelligence (GenAI) technologies continue to advance, GenAI-based items are expected to play a significant role in the item development process within educational and professional testing programs. In this study, GenAI-based items refer to items initially generated by large language models (LLMs) using a predefined prompting strategy and subsequently undergo human review prior to administration. Prior studies have examined AI-generated items. Concerns regarding their bias, clarity and overall quality have been addressed in literature and have challenged their broader use. Empirical evidence regarding GenAI-based items that incorporate human review and refinement under actual testing conditions remains limited. Such investigation is particularly important in high-stakes testing, where item quality directly affects score reliability and test outcome.

2 Aims

The present study seeks to evaluate GenAI-based multiple-choice questions (MCQs) developed for medical assessments in a high-stakes testing context. The study was guided by two research questions: (1) Are these items comparable to human-developed items in terms of content quality? and (2) Do they exhibit psychometric characteristics comparable to those of human-developed items? The evaluation focused on item content quality and key psychometric properties, including the quality of item stems and item response options, classical item statistics, item response theory (IRT)-based item parameter estimates, and item flagging rates.

3 Methods

The study integrates pre-administration content review with post-administration psychometric analyses to provide a comprehensive evaluation of item quality and the comparability of the GenAI-based and subject matter experts (SME)-developed items.

3.1 Item Development

The GenAI-based items examined in this study were developed in 2025 using GPT-4.1. A typical MCQ consists of an item stem and a set of response options. In medical assessments, item stems typically present clinical scenarios that incorporate patient symptoms, medical history, physical examination findings, and/or laboratory results. The prompting strategy was evaluated in prior studies (Walker et al., 2025A, 2025B; Makinde, 2025) and includes the following components: 1). question stem requirements following established medical examination pattern and difficulty; 2). clinical scenario templates containing patient demographics, presenting symptoms, examination findings, and laboratory results; 3). five-option

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 275–281
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

multiple-choice format with detailed rationales for each option for SME review; 4). relevant organization style rules; 5). chain-of-thought reasoning to ensure clinical coherence; 6). a three-item quiz format for each clinical topic to produce item drafts that assessed the same topic in different ways.

3.2 Item Content Evaluation

The evaluation included a systematic content review and revision process conducted by six experienced item writers (i.e., SMEs) and two Exam Content Developers. For each clinical topic, the SMEs panel jointly reviewed GenAI-based items until selecting one to use (typically in the first set of three items). Necessary revisions were made with the help of the Exam Development team members and tracked.

Afterwards, the number and types of changes were categorized to assess item content quality. The revisions were assigned to one of seven categories – laboratory value formatting, unit formatting, text changes for style and clarity, distractor revision/replacement, terminology modification, clinical focus, and patient details.

3.3 Psychometric Analysis

3.3.1 Item response options

Response options were evaluated at both the item and test-section levels, with particular attention to distractor functioning. Distractors were classified as functioning or non-functioning using two criteria (Ebel & Frisbie,1991; Lakens, 2017; Loudon & Macias-Muñoz, 2018; Osterlind, 1998). First is selection proportion (P_{ij}). A distractor j in item i was classified as a non-functioning distractor (NFD) if it was selected by fewer than 5% of examinees, i.e., $P_{ij} = \frac{n_{ij}}{N} < 0.05$, where N denotes the number of examinees responding to item i and n_{ij} denotes the number selecting distractor j on item i. Second is discrimination index D_j. Distractor discrimination was evaluated by comparing selection rates between upper- and lower-ability groups. A distractor was classified as non-functioning when its selection rate was higher among the upper-ability group than among the lower-ability group. That is, $D_j = P_{Uj} - P_{Lj} > 0$, where $P_{Uj} = \frac{n_{Uj}}{N_U}$ and $P_{Lj} = \frac{n_{Lj}}{N_L}$.

3.3.2 Statistical flagging rate:

Statistical flagging was based on classical item statistics, including item percent correct, point-biserial correlation, or response time. Although specific criteria may vary across testing programs, common flagging rules include low or negative point-biserial correlation, extreme item difficulty or response time, and combinations of these indicators. In the present study, the flagging criteria used in the operational review process for a certification examination were applied consistently to both SME-developed and AI-assisted items. The resulting flagging rates were then compared to assessing the overall psychometric comparability of the two item types.

3.3.3 Item Statistics and Item Parameters

Item statistics and model-based parameters are central to evaluating item functioning and informing test scoring. Accordingly, examining these characteristics across item types provides important evidence regarding the feasibility of using GenAI-based items. Specifically, the distribution of response time, percent correct (i.e., classical item difficulty), point-biserial correlation coefficients, and IRT item difficulty were compared between the two item types.

4 Results

4.1 Data

The study sample included 3,514 examinees who completed a high-stake medical certification examination during the first half of 2026. The examination was administered as a timed, computer-based test consisting of 300 MCQs designed to assess medical and surgical knowledge and skills. Both SME-developed and GenAI-based pretest items were embedded within the examination. Four test forms were administered, with each form associated with three 50-item pretest sections, yielding 12 unique pretest sections in total. Examinees were randomly assigned to one of these pretest sections, and each pretest item was administered to approximately 282 to 305 examinees).

4.2 Item Content

Sixty-two GenAI-based items were reviewed and revised prior to administration. The 62 items contained 50 items for a complete pretest section and the remaining 12 were retained as backup

items in case content issues were identified during Forms Review (although none were subsequently chosen for replacement by two pairs of SMEs who were blinded to the AI assisted nature of the items). Content revisions were classified into seven categories described previously. Across categories (Appendix Table 1), the number of changes per item ranged from 0 to 18, with the average number of changes per item ranging from 0.27 (less than one change) to 9.48 (more than 9 changes). Revisions were most frequent for stem style or clarity, distractor content, and patient details (Categories 3, 4, and 7), with approximately 85% to 98% of items receiving revisions in these areas. In contrast, relatively few revisions were made to clinical focus and lab formatting (Categories 1 and 6).

To account for potential variation across content areas, GenAI-based and SME-developed items were compared within the same content blueprint and clinical topic. Fifteen matched pairs were identified based on their clinical topics. GenAI-based items required fewer revisions than SME-developed items across most categories, with the exception of distractor revision or replacement (Category 4). For both item types, revisions remained most frequent in Categories 3, 4, and 7 (Figure 1). Overall, the revision pattern for GenAI-based items after matching content was consistent with that observed in the unadjusted condition.

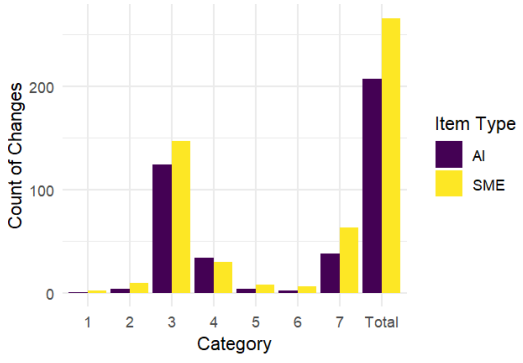


Figure 1: Count of changes across 7 categories for both item types

4.3 Psychometric analyses

Item quality was further evaluated using psychometric analyses, with SME-developed items serving as the reference group. Item analysis was conducted in R to obtain response time (RT), percent correct (p+), and point-biserial correlation (r_{pb}). A unidimensional Rasch model was implemented in the computer software WINSTEPS to estimate item difficulty and person ability for both item types. A concurrent calibration was carried out to ensure model parameters at different test forms are on the same scale. Mean ability estimates ranged from 1.68 to 1.76, with standard deviation (SD) ranging from 0.44 to 0.55. Prior to comparing item statistics across the 12 pretest sections (4 forms x 3 pretest sections), mean ability was examined for the corresponding examinee groups. Equivalence across groups was assessed using 'two one-sided t-tests' (TOST) procedure with an equivalence margin of +/- 0.15 logits. All pairwise comparisons satisfied the equivalence criterion ($p < .05$), supporting the comparability of examinee ability across forms and pretest sections.

4.3.1 Item Statistics

Pretest sections were constructed to achieve balanced content coverage across test forms. Response data was obtained from 12 unique sections. Section 1 contained 50 AI-assisted pretest items, whereas Sections 2-12 contained SME-developed items. Given the aligned ability distributions across examinee groups and balanced content coverage, consistent item performance was expected across item types. As shown in Table 1, mean response time (RT), percent correct (P+), and point-biserial correlation (r_{pb}) are generally similar across sections. The largest differences were 7 seconds in RT, 0.08 in P, and 0.04 in r_{pb}.

Section	N	RT	P+	r_{pb}	b
1	297	69.02	0.72	0.12	0.28
2	287	69.01	0.69	0.14	0.57
3	302	65.42	0.72	0.12	0.21
4	299	66.85	0.74	0.11	0.25
5	305	66.15	0.70	0.12	0.42
6	286	64.17	0.69	0.14	0.41
7	282	63.02	0.70	0.11	0.48
8	294	62.73	0.73	0.10	0.20
9	292	65.86	0.72	0.12	0.28
10	289	61.85	0.74	0.11	0.01
11	297	66.04	0.66	0.13	0.69
12	284	67.03	0.68	0.12	0.39

Note: N = Number of examinees in each section

Table 1: Average Item Statistics across Pretest Sections

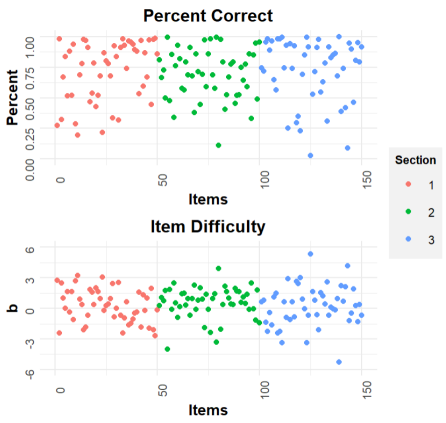


Figure 2: Percent correct and b-parameter estimates for items in Sections 1-3

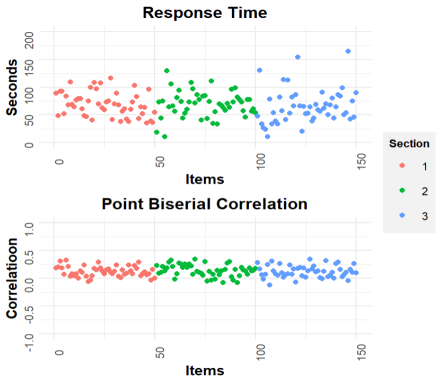


Figure 3: Response time and point-biserial correlation for items in Sections 1-3

Response time was slightly longer in Section 1 (AI-assisted) than in other sections, possibly reflecting its higher stem and response option word count. Median stem and response option lengths were 117 and 21 words, respectively, compared with approximately 105 (stem) and 18 (response option) words in most other sections (Appendix Table 2).

Figure 2 and 3 present the distributions of item statistics and Rasch b-parameter estimates for Sections 1–3, which were administered with the same operational test form. Section 1 contained AI-assisted items, whereas Sections 2 and 3 contained SME-developed items, allowing their distributional patterns to be examined under the same operation form. The distributions appeared visually similar across sections in terms of range, density, and spread. Section 2 showed a slightly higher mean b-parameter estimate, likely influenced by a greater concentration of more difficult items.

Given the importance of Rasch item difficulty b-parameters in test scoring, additional analyses were conducted to assess their distributions for AI-assisted and SME-developed items. Wasserstein two-sample tests (Ramdas et al., 2017; Panaretos & Zemel, 2019) were performed at both the section and item-type levels. Section-level comparisons contrasted Section 1 (GenAI-based items) with each of Sections 2–12 (SME-developed items); all tests were nonsignificant (p=.1622 to .9256). At the item-type level, the Wasserstein distance was 0.2608 (p=.6838). Collectively, these analyses detected no statistically significant differences in the distributions of b-parameter estimates between the two item types. Consistent with these results, the Q-Q plot (Figure 4) showed quantiles closely aligned with the identity line, providing visual support for similar distributional patterns.

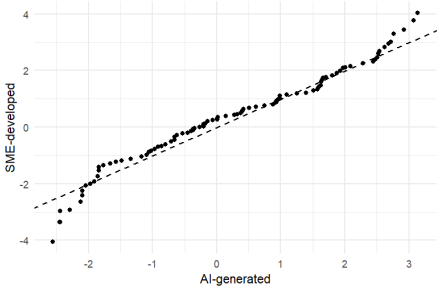


Figure 4: Q-Q plot of b-parameter estimates for GenAI-based and SME-developed items

4.3.2 Statistical Flag Rate

Four operational flagging criteria were applied to evaluate pretest item statistics. As shown in Figure 5, AI-assisted items had relatively high flag

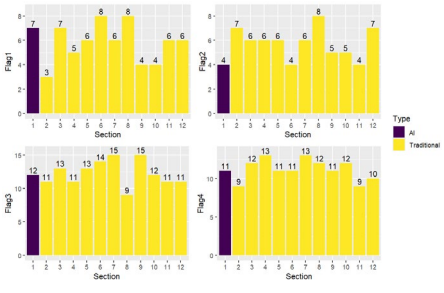


Figure 5: Number of flags across sections based on 4 rules

counts under criterion 1 (i.e., hard item) and low counts under criterion 2 (low point-biserial correlation), but no unusual pattern evident across the four criteria overall. We also examined the number of items flagged within each section. The GenAI section had 21 flagged items (42%), the highest rate observed, whereas the SME-developed sections ranged from 14 to 20 items (28%-40%). Although the GenAI section showed the highest number of flagged items, several SME-developed sections had similar counts (Appendix Table 3).

4.3.3 Item response options

Response data on item distractors from all 600 pretest items were analyzed. During content review, some item response options were removed, resulting in a small number of four-option items (Appendix Table 4). The final analysis therefore included 2,384 distractors rather than 2,400. To compute the discrimination index, examinees at the top 33% of the ability distribution were classified as the upper group, and those in the bottom 33% as the lower group. Using the selection-frequency criteria, Table 2 showed that the proportion of functioning distractors (FD) across SME-developed sections ranged from 34% to 46%, with most sections near 37%. Section 1, which contained GenAI-based items, had an FD rate of 36%, close to the median of the SME-developed sections. Under the discrimination-index criterion, FD rate ranged from 76% to 84%, substantially higher than those based on selection frequency, with the GenAI section at the lower end of this range. Pearson's Chi-Squared test was conducted to examine the association between item type (GenAI vs. SME) and distractor functioning. No statistically significant association was found under either criterion, x^2 (1) = 0.19316, p = .6603 for selection frequency, or x^2 (1) = 2.5542, p = .110 for discrimination index.

Section	S.F.	D. I.	Both
1	0.36	0.76	0.28
2	0.43	0.81	0.31
3	0.34	0.84	0.23
4	0.35	0.84	0.27
5	0.36	0.76	0.25
6	0.36	0.82	0.29
7	0.37	0.82	0.30
8	0.39	0.83	0.31
9	0.38	0.80	0.29
10	0.35	0.78	0.25
11	0.46	0.78	0.33
12	0.37	0.79	0.27
Total	0.38	0.80	0.28

Note: S.F. = Selection frequency criterion
D.I. = Discrimination index criterion
Both = both criteria are considered

Table 2: Functioning Distractor Rate across Sections

5 Conclusions

This study evaluated the quality of GenAI-based items using content-review data and examinee responses from a high-stakes medical certification examination. Under content-matched conditions, GenAI-based items required fewer revisions than SME-developed items, suggesting comparable content quality and potential gains in item-development efficiency.

Psychometric analyses showed that GenAI-based items performed comparably to SME-developed items across classical item statistics, Rasch b-parameter estimates, distractor functioning, and operational flagging indicators. Collectively, these findings provide empirical support for the feasibility of incorporating GenAI-based items into operational assessment development.

Both GenAI-based and SME-developed items underwent rigorous human review and revision before administration to meet the requirement for operational use; therefore, the GenAI-based items evaluated in this study reflect a combined AI–human development process rather than fully automated item generation.

Several limitations should be noted. First, the study included 50 GenAI-based items, and replication with larger and more diverse item samples is needed to strengthen generalizability. Second, the Wasserstein analyses assessed whether distributional differences could be detected but did not formally establish distributional equivalence. Future research should apply explicit similarity or equivalence frameworks. Third, item parameters were examined only under a Rasch model. Future studies should evaluate item performance under additional IRT models, including the two- and

three-parameter logistic models, to extend the relevance of the findings to testing programs that use these models.

6 References

Ebel, R. L. & Frisbie, D.A. (1991). *Essentials of educational measurement*. 5th edition. Englewood Cliffs, N.J.: Prentice Hall, Englewood Cliffs.

Lakens, D. (2017). *Equivalence Tests: A Practical Primer for t Tests, Correlations, and Meta-Analyses.* Social Psychological and Personality Science, 8(4), 355–362. https://doi.org/10.1177/1948550617697177

Loudon, C. & Macias-Muñoz, A. (2018). *Item statistics derived from three-option versions of multiple-choice questions are usually as robust as four- or five-option versions: implications for exam design*. Advances in Physiology Education 42 (4): 565–575. https://doi:10.1152/advan.00186.2016

Makinde, H., Walker, M., & Dallas, A. (2025, April 23-26). AI Generation of Medical specialty Items with LLM. Annual Meeting of the National Council on Measurement in Education Conference, Denver, CO, United States.

Osterlind, S. J. (1998). *Constructing test items: Multiple-choice, constructed- response, performance, and other formats* (2nd ed.). Kluwer Academic Publishers.

Panaretos, V. M., & Zemel, Y. (2019). Statistical aspects of Wasserstein distances. *Annual Review of Statistics and Its Application, 6*, 405–431. DOI: 10.1146/annurev-statistics-030718-104938.

Ramdas, A., García Trillos, N., & Cuturi, M. (2017). *On Wasserstein two-sample testing and related families of nonparametric tests*. Entropy, *19*(2), 47. DOI: 10.3390/e19020047

Schuirmann, D. J. (1987). *A comparison of the two one-sided tests procedure and the power approach for assessing the equivalence of average bioavailability.* Journal of Pharmacokinetics and Biopharmaceutics, 15(6), 657–680. https://doi.org/10.1007/BF01068419

Walker, M., Richburg, A., Makinde, H., Dallas, A., Jiang, Y., & Gregg, N. (2025, October 27-29). *SME comparison of PubMed RAG vs non-RAG items and fact-checking.* [Paper presentation]. Annual Meeting of the Artificial Intelligence in Measurement and Education Conference, Pittsburgh, PA, United States.

Walker, M., Makinde, H., Richburg, A., Dallas, A., & Goodman, J. (2025, October 27-29). *From theory to practice, high-stakes item development and review with Generative AI.* [Paper presentation]. Annual Meeting of the Artificial Intelligence in Measurement and Education Conference, Pittsburgh, PA, United States

A Appendices

	Cat. 1	**Cat. 2**	**Cat. 3**	**Cat. 4**
Items with a change	19.4%	38.7%	98.4%	87.1%
N. changes per item	0.27	0.56	9.48	2.35
Minimum changes	0	0	0	0
Maximum changes	3	4	18	5
	Cat. 5	**Cat. 6**	**Cat. 7**	
Items with a change	21.0%	12.9%	85.5%	
N. changes per item	0.34	0.13	2.63	
Minimum changes	0	0	0	
Maximum changes	4	1	8	

Note: Category 1: Lab formatting; Category 2: Unit formatting; Category 3: Text changes for style and/or clarity; Category 4: Distractor content (revising/replacing); Category 5: Medical terminology (revising or replacing); Category 6: Clinical focus (changing the lead-in to meet the assigned task); Category 7: Patient details:

Appendix Table 1: Summary of Item Revisions in GenAI-based Items

Section	Stem	Option
1	117	21
2	114	20
3	114	21
4	93.5	20
5	104	17.5
6	102	17
7	104	14
8	104	17.5
9	112	17
10	109	16.5
11	108	19.5
12	106	18

Appendix Table 2: Median of Word Counts per Item in Each Section

Section	N	N Flagged Items	% Flagged items
1	50	21	0.42
2	50	14	0.28
3	50	17	0.34
4	50	18	0.36
5	50	18	0.36
6	50	20	0.4
7	50	20	0.4
8	50	19	0.38
9	50	18	0.36
10	50	16	0.32
11	50	17	0.34
12	50	14	0.28

Appendix Table 4: Frequency and % of Flagged Items across Sections

Section	N. items (4-option)	N. items (5-option)	N. Distractors
1	1	49	199
2	3	47	197
3	0	50	200
4	1	49	199
5	2	48	198
6	1	49	199
7	2	48	198
8	2	48	198
9	1	49	199
10	0	50	200
11	2	48	198
12	1	49	199
Total	16	584	2384

Appendix Table 3: Frequency of Items and Distractors across Sections

When One Benchmark Hides Many Truths: Mixture-IRT and LLM Difficulty Prediction

Enis Dogan[1], Burhan Ogut[2]
[1]National Center for Education Statistics [2]American Institutes for Research

Abstract

Studies that validate LLM-predicted item difficulty conventionally benchmark predictions against a single set of item parameter estimates fitted to the full examinee sample, treating the ordering of item difficulty as homogeneous across the population. We test that assumption using responses from 14,501 examinees to a 29-item national university-entrance mathematics examination, together with difficulty predictions from a locally hosted multimodal LLM (gemma4:26b) elicited through blind, position-counterbalanced pairwise comparison and aggregated via a Bradley–Terry model. A three-class mixture Rasch model shows latent classes that differ in proficiency, intended collegiate field, and instructional track. LLM predictions align most closely with the largest, highest-proficiency class (ρ = .69) and markedly less closely with the two lower-proficiency classes (ρ = .48 and .49), while the conventional single-population benchmark (ρ = .63) sits between these values and conceals the spread entirely. Because a judgment that one item is harder than another presumes a solution process (Mislevy & Verhelst, 1990), these predictions may encode an examinee model that is broadly correct without being universal. We recommend that LLM difficulty-prediction studies examine whether latent classes exist in their validation samples and, where they do, report alignment separately by latent class.

1 Introduction

A growing literature uses machine learning and large language models (LLMs) to predict the empirical difficulty of test items, supporting item development, pre-calibration, and assessment design. Several studies presented at 2025 AIME-Con report encouraging results (e.g., Mojoyinola et al., 2025; Adegoke et al., 2025; Han et al., 2025; Li et al., 2025), as does a growing wider literature (AlKhuzaey et al., 2024; Kolesnikova et al., 2026).

The prevailing validation convention in this literature benchmarks predictions against a single set of empirical difficulty estimates, obtained by fitting one item response theory (IRT) model to a full examinee sample and treated as the "ground truth". This convention embeds a strong and rarely examined assumption that the ordering of item difficulty is homogeneous across the population.

If the examinee population comprises qualitatively distinct latent subgroups for whom the item-difficulty orderings differ, consistent with subgroups employing different solution strategies, a benchmark computed against a single set of item parameter estimates averages over that heterogeneity. Predictions may then align well with the difficulty estimates of some subgroups and poorly with those of others, and such a benchmark will both understate the best attainable alignment and obscure systematic misalignment for particular subpopulations.

This heterogeneity has particular significance for AI-based difficulty prediction. A judgment that one item is harder than another implies something about the work a test-taker must do to answer it, and so presumes a solution process. Yet, test-takers employ a variety of strategies in responding to items. For example, Mislevy and Verhelst (1990) constructed an example in which items relatively hard under one solution strategy are relatively easy under another. If an LLM's predictions encode a particular model of how examinees approach items, that model may be broadly correct without being universal, aligning with the subpopulations whose approach it resembles and diverging from those it does not.

Using item responses from a national university-entrance mathematics examination (N = 14,501) and difficulty predictions from a multimodal LLM elicited through blind pairwise comparison, we ask: (1) Does a multi-class latent structure exist in a real high-stakes dataset, and is it substantively interpretable? (2) Do the item parameters differ enough across classes to matter for differential validation? (3) Does the alignment between LLM predictions and empirical difficulty vary by latent class, and if so, which class does the LLM most closely track? We find that alignment varies substantially across classes and that the aggregate benchmark conceals this variation.

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 282–288
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

2 Related Work

2.1 Predicting Item Difficulty from Item Content

Model-based prediction of item difficulty from item content predates the current generation of language models. Early approaches relied on readability indices, item metadata, and hand-coded cognitive features, whereas contemporary work replaces or augments these with *learned* text representations. For instance, Han et al. (2025) fine-tuned four transformer architectures on Grade 6 ELA items, reaching correlations near .70 with empirical difficulty. Li et al. (2025) compared fine-tuned small language models against larger general-purpose LLMs on medical licensure items and found the smaller fine-tuned models often stronger, and related work extends this line to international large-scale assessment (Mojoyinola et al., 2025; Adegoke et al., 2025). Across this literature, the reported outcome is a benchmark computed against a single set of item parameter estimates fitted to the full sample.

2.2 Pairwise Comparison as an Elicitation Format

A parallel line of work elicits difficulty judgments comparatively rather than absolutely, drawing on the comparative-judgment tradition in educational measurement (Bramley, 2010; Crompvoets et al., 2020). Kolesnikova et al. (2026) provide the most direct precedent, comparing absolute judgment against pairwise comparison across six domains of primary-school mathematics and aggregating pairwise judgments via the Bradley–Terry model (Bradley & Terry, 1952). They found pairwise elicitation generally produced stronger alignment with empirical difficulty, and that few-shot prompting conferred little benefit. Their design randomized item order within each pair to mitigate order effects, though the magnitude of those effects was not reported. The present study adopts pairwise elicitation and additionally counterbalances rather than randomizes order, which permits the position-consistency diagnostic in Section 4.3.

2.3 Population Heterogeneity in Item Difficulty

Mixture item response models relax the assumption that a single set of item parameters can be obtained for an entire examinee population. Rost (1990) integrated latent class analysis with the Rasch model, allowing distinct classes each with its own item-difficulty vector. Mislevy and Verhelst (1990) framed the same idea substantively, modeling responses when examinees employ different solution strategies. Bolt et al. (2002) applied mixture Rasch models to test speededness, showing that latent classes can capture interpretable differences in how subgroups engage with items rather than mere noise.

2.4 The Gap

These two literatures have not been brought into contact. Difficulty-prediction studies validate against a population-aggregate benchmark, but the mixture-IRT research establishes that such a benchmark can average over substantively distinct subpopulations. If distinct latent classes do exist, benchmarking predictions against a single set of item parameter estimates would be akin to a weighted average of subgroup-specific alignments that may differ considerably. No existing study, to our knowledge, reports LLM-difficulty alignment separately by latent class.

3 Methods

3.1 Analytic Logic

Our analysis proceeds in three linked stages. First, we establish a single-population ("one-class") Rasch difficulty ordering, replicating the convention used throughout the LLM difficulty-prediction literature. Second, we relax the homogeneity assumption using mixture Rasch modeling. Model selection and covariate profiling assess whether the resulting classes are statistically supported and substantively interpretable rather than merely fit artifacts. Third, we generate an independent, LLM-derived difficulty ordering and evaluate its agreement with both the aggregate benchmark and the class-specific difficulty estimates.

3.2 Instrument

Data are item-level responses and the original item content from the mathematics subtest of a 2004 national university-entrance examination administered in an OECD-member country. The original mathematics subtest comprised 44 items covering numbers and operations, algebra, word problems, and geometry. Nine trailing items with 20% or more not-reached responses were

removed to address speededness, leaving 35. One item with a proportion-correct of .98 was removed for lack of discriminating information, yielding 34. An exploratory factor analysis (tetrachoric correlations, promax rotation) identified five coordinate-geometry items loading on a secondary factor; these too were removed, yielding a content-coherent set of 29 items suitable for unidimensional IRT modeling.

3.3 Sample

The analytic sample (N = 14,501) was 43.2% female and 56.8% male. By instructional track, 61.3% of students were enrolled in a mathematics–science program, 33.3% in a mathematics–social studies program, and 5.4% in a social-sciences program. Examinees also indicated the collegiate field they intended to apply to: a large majority (68.2%) reported interest in a STEM-related field, while the remainder indicated fields that admit students based on an equally weighted mathematics–social studies exam score.

3.4 Mixture Rasch Modeling

We fit a one-class Rasch model as a baseline and replication of prior literature, then two-, three-, and four-class mixture Rasch models using the *psychomix* package in R (raschmix, saturated score model; Frick et al., 2012).

In preparing the dataset for the IRT estimation, skipped items — those left blank before the examinee's final recorded response — were treated as incorrect, on the assumption that the examinee saw the item and elected not to answer. Not-reached items, those following the last item attempted, were treated as not-presented rather than incorrect throughout estimation, since coding them incorrect could produce a spurious speededness class rather than a substantive ability class.

Each multi-class solution was estimated from 200 independent random starts executed in parallel with an EM iteration cap of 500, retaining the highest-log-likelihood solution to guard against local maxima. Model enumeration considered AIC, BIC, marginal fit improvement, convergence behavior across starts, classification certainty, and substantive interpretability. Class membership was profiled against intended collegiate field, instructional track, and sex, and against response accuracy, defined as the percentage of attempted items answered correctly and the number of items not reached.

3.5 LLM Difficulty Prediction

Item difficulty was independently estimated via a blind, stateless pairwise-comparison pipeline using a local multimodal LLM (gemma4:26b, via Ollama). Each item was represented as an image crop of the original question, preserving the original text, mathematical notation, and any figures, so that the model judged items in their native visual form rather than a transcribed rendering. Two items included figures and one included a table.

Difficulty predictions were elicited as a complete round-robin of pairwise comparisons among the 29 analytic items. Each unordered pair was presented twice, once in each presentation order, at temperature 0, yielding 812 independent judgments (406 unordered pairs × 2 orders). Each call presented two item images labeled ITEM A and ITEM B and required a forced binary verdict (ties disallowed) with a one-clause rationale, returned as constrained JSON. The model's chain-of-thought reasoning mode was disabled, so each response consisted of the verdict and rationale alone. Calls carried no conversation history, so each judgment is independent of every other and position bias can be assessed directly. These 812 judgments form the basis for all analyses reported in Section 4.

Judgments were aggregated into a continuous latent difficulty scale using the Bradley–Terry model (Bradley & Terry, 1952), under which the probability that item i is judged harder than item j is $\sigma(\theta_i - \theta_j)$, where θ is the item's difficulty on a log-odds scale and σ is the logistic function. Parameters were estimated by maximum likelihood over the 812 comparisons, with the scale identified by fixing one item's θ at zero. One item lost every comparison in which it appeared, for which no finite maximum-likelihood estimate exists; a small ridge penalty ($\lambda = 0.05$) on θ was applied to keep all estimates finite. The resulting scale was compared with empirical difficulty from the one-class baseline and each class of the three-class solution using Spearman's ρ and Pearson's r.

4 Results

4.1 Model Selection

Estimation was stable for the one-, two-, and three-class solutions. The three-class solution

was retained based on convergence behavior, classification quality, decelerating marginal fit improvement, and the substantive coherence of the resulting classes (Section 4.2). Fit improved monotonically across the one-, two-, and three-class solutions. BIC fell from 406,064.00 at one class to 401,661.43 at two, a gain of 4,403 points, and to 400,006.09 at three, a gain of 1,655. Classification quality was adequate and stable: mean maximum posterior probability was .907 at two classes and .857 at three, with relative entropy of .688 and .687 respectively, indicating that the third class did not degrade the certainty with which examinees are assigned.

The four-class solution was sought but eventually abandoned because it showed estimation instability and produced degenerate estimates. Although the additional class had a mixing proportion of 16.29%, multiple score-model parameters were estimated at $-\infty$, indicating that the solution was not identifiable in a substantively useful way.

4.2 Class Characterization

The three latent classes formed a clear, monotone proficiency gradient. Because not-reached responses were treated as missing, classes are distinguished primarily by response accuracy (correctness) on attempted items rather than by number of items attempted. Class 2 was the largest and highest-performing (n = 8,887; 61.3%; 73.7% correct on attempted items), Class 1 was intermediate in proficiency (n = 3,648; 25.2%; 46.3% correct on attempted items), and Class 3 was smallest and lowest-performing (n = 1,966; 13.6%; 35.4% correct on attempted items). Against a chance rate of roughly 20% for five-option items, even the lowest class performed above chance, so the gradient is one of proficiency rather than a contrast between engaged and guessing examinees. Relative entropy was .687, with a mean maximum posterior class probability of .857. Table 1 reports the covariate profile.

	Class 2	Class 1	Class 3
% of sample	61.3	25.2	13.6
% correct (attempted items)	73.7	46.3	35.4
Mean items not reached	0.03	0.85	2.61
Collegiate field (%): STEM	92.5	35.7	18.4
Female (%)	36.7	61.8	38.3
Track (%):			
Mathematics-Science	79.6	36.5	24.7
Mathematics-Social Studies	20.1	56.2	50.6
Social Studies	0.3	7.4	24.7

Table 1. Covariate profile by latent class, ordered by proficiency.

The intended collegiate field separated the classes most sharply ($V = .67$), with STEM field at 92.5% in Class 2, against 35.7% in Class 1 and 18.4% in Class 3, $\chi^2(2) = 6{,}442.8$, $p < .001$. Examinees also differed by instructional track by latent class ($V = .39$), with mathematics–science enrollment falling from 79.6% in Class 2 to 36.5% in Class 1 and 24.7% in Class 3, while social-studies enrollment rose from 0.3% to 7.4% to 24.7%, $\chi^2(4) = 4{,}306.6$, $p < .001$.

Sex was significantly but far more weakly associated with class membership ($V = .22$) and, importantly, was not monotone in proficiency: the intermediate Class 1 was majority female (61.8%), whereas both the highest-performing Class 2 (36.7% female) and the lowest-performing Class 3 (38.3% female) were majority male, $\chi^2(2) = 689.1$, $p < .001$.

Non-response followed the proficiency gradient. Class 2 examinees reached essentially the whole set (mean 0.03 items not reached), Class 1 substantially less so (mean 0.85 items not reached), and Class 3 least (mean 2.61 items not reached). The two lower-proficiency classes thus differed from the highest not only in number of correct answers but in test-taking endurance under time pressure, consistent with the speededness mechanism documented by Bolt et al. (2002).

4.3 Pairwise Comparisons and Bradley–Terry Scale

Two diagnostics bear on the quality of the resulting scale. The first is internal fit: how well the fitted Bradley–Terry scale reproduces the judgments from which it was estimated. The scale correctly predicted 750 of the 812 judgments (92.36%), indicating that the LLM's judgments were well summarized by a single latent difficulty dimension.

The second diagnostic is comparison accuracy: for each of the 812 judgments, whether the item the LLM selected as harder is in fact the harder of the pair under a given empirical benchmark. Unlike internal fit, this is a validity measure and varies by benchmark; it is reported alongside the correlations in Section 4.4. Because it is computed judgment by judgment rather than from the aggregated scale, it provides a check on alignment that does not depend on the Bradley–Terry aggregation or on the choice of correlation coefficient.

Position consistency is a third property of the judgment set. Of the 406 unordered pairs tested in both presentation orders, the LLM's judgment reversed with position in 54 pairs (13.3%) (judging item A as more difficult when A is presented first, then B as more difficult when B is presented first). Order sensitivity of this kind is a documented property of LLM judges rather than a peculiarity of this pipeline: position bias is a recognized failure mode in the LLM-as-a-judge literature (Zheng et al., 2023; Shi et al., 2024). Reversals here were not evenly distributed across items. The two most reversal-prone items appeared in 11 and 10 reversed pairs respectively, together accounting for 19.4% of all reversals, while five of the 29 items appeared in one or zero reversals. The most reversal-prone item had an empirical difficulty near the center of the scale.

4.4 Differential LLM Alignment Across Classes

Figures 1 and 2 report the agreement between the LLM's Bradley–Terry scale and empirical difficulty (β) from the single-population benchmark and each latent class: Figure 1 gives rank correlations, and Figure 2 gives comparison accuracy computed directly from the 812 individual judgments, where a comparison counts as accurate if the item the LLM judged harder is the one with the higher empirical difficulty estimate under that benchmark.

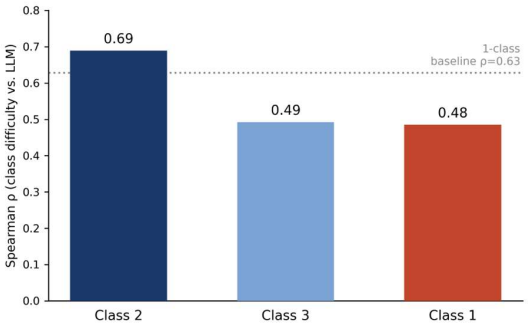


Figure 1. Spearman correlation between empirical item difficulty and LLM-predicted difficulty, by latent class. The dotted line shows the single-benchmark (one-class) correlation.

The LLM aligns most closely with the highest-proficiency class (Class 2: ρ = .69, r = .63, 72.91% accuracy in pairwise comparisons) and least closely with the intermediate and lowest-proficiency classes (Class 1: ρ = .48, r = .45, 65.02% accuracy; Class 3: ρ = .49, r = .52, 65.27% accuracy), a spread of .20 in rank correlation, corroborated by both the Pearson correlation and a roughly eight-percentage-point gap in comparison accuracy. Because accuracy is computed from individual pairwise judgments rather than the aggregated scale, the pattern is not an artifact of the correlation metric or of Bradley–Terry aggregation.

Critically, the single-population benchmark (ρ = .63, r = .57, 69.95% accuracy) — the figure a conventional validation would report — sits between these values on both metrics, concealing that alignment with the two lower-proficiency classes, together 38.8% of the population, is just over two-thirds as strong by rank correlation as alignment with the highest-performing class. Against a chance rate of 50% for forced binary comparisons, the accuracy gap is a larger share of the attainable signal than it first appears: the LLM's judgments are correct 22.9 percentage points above chance for the highest-proficiency class but only about 15 points above chance for the two lower classes, roughly a third less. This is the central empirical finding motivating our argument against single-benchmark validation.

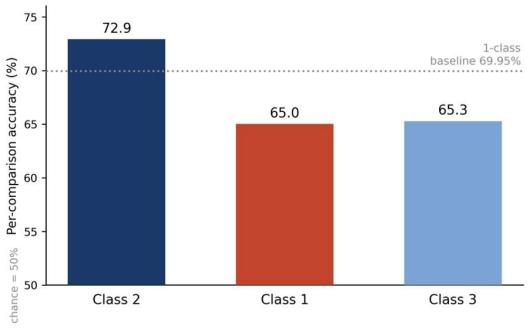


Figure 2. Comparison accuracy of LLM difficulty judgments against empirical item difficulty, by latent class (ordered by proficiency). The dotted line shows the single-benchmark (one-class) accuracy. The axis begins at the 50% chance rate for forced binary comparisons.

One nuance the three-class solution makes visible is that alignment is not monotone in proficiency. The intermediate class (ρ = .48) is tracked marginally less well than the lowest class (ρ = .49), and the two are indistinguishable at this item count. The finding is best characterized as the LLM tracking the highest-proficiency majority class well and both lower-proficiency classes comparably less well, rather than as a smooth decline across the ability range.

5 Conclusions and Limitations

This study showed that a single, population-aggregate correlation between LLM-predicted and empirical item difficulty can mask substantial heterogeneity across an examinee population's latent subgroups. Using a three-class mixture Rasch model of a national mathematics examination, we found that a locally hosted multimodal LLM's pairwise judgments aligned strongly with the highest-proficiency class (ρ = .69) but considerably less well with the intermediate and lowest-proficiency classes (ρ = .48 and .49), a gap the conventional benchmark (ρ = .63) obscures. The pattern held at the level of both aggregated rank correlation and individual pairwise judgments, and the classes were externally corroborated by intended collegiate field, instructional track, and not-reached behavior.

One plausible account for the differential alignment is that difficulty for high-proficiency examinees is closely tied to intrinsic item complexity — solution steps, abstraction, conceptual depth an expert-style judge would readily recognize — while difficulty for lower-proficiency examinees is shaped by factors an LLM cannot access from item text or image alone: guessing, differential content exposure and other factors. A second, compatible account follows from Mislevy and Verhelst (1990): if the classes differ partly in solution strategy, then perhaps the LLM's predictions encode a particular model of how an examinee approaches an item, valid for the subpopulation whose approach it resembles.

The practical implication is clear. A single population-level benchmark cannot establish that a predictor is equally trustworthy for every subgroup a testing program serves. We recommend that LLM difficulty-prediction studies examine whether latent classes exist in their validation samples and, where they do, report alignment separately by latent class. Several limitations qualify these conclusions:

- **Item and sample exclusion criteria.** The item reduction and the removal of examinees with extreme scores or excessive missingness were principled but potentially consequential; since not-reached responses and extreme scores may relate to latent-class formation, these decisions could shape the composition of the lowest-proficiency class specifically. No formal sensitivity analysis varying these thresholds was conducted.
- **Single LLM, single pass.** Findings rest on a single LLM at one temperature in a single pass; the 13.3% reversal rate indicates instability even within that pass, and generalization to other LLMs is untested.
- **Single instrument and population.** All findings come from one mathematics subtest. The results may not generalize elsewhere.

Future work should extend this analysis to additional LLMs, instruments, and populations, and should pair class-specific alignment reporting with the identifiability diagnostics this study found necessary.

References

Adegoke, H. O., Du, Y., & Dwyer, A. (2025). Medical item difficulty prediction using machine learning. In *Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con), Works in Progress* (pp. 185–190). National Council on Measurement in Education.

AlKhuzaey, S., Grasso, F., Payne, T. R., & Tamma, V. (2024). Text-based question difficulty prediction: A systematic review of automatic approaches. *International Journal of Artificial Intelligence in Education*, 34(3), 862–914.

Bolt, D. M., Cohen, A. S., & Wollack, J. A. (2002). Item parameter estimation under conditions of test speededness: Application of a mixture Rasch model with ordinal constraints. *Journal of Educational Measurement*, 39(4), 331–348.

Bradley, R. A., & Terry, M. E. (1952). Rank analysis of incomplete block designs: I. The method of paired comparisons. *Biometrika*, *39*(3/4), 324–345.

Bramley, T. (2010, June). Locating objects on a latent trait using Rasch analysis of experts' judgments. Paper presented at the conference "Probabilistic Models for Measurement in Education, Psychology, Social Science and Health," Copenhagen, Denmark.

Bramley, T. (2010, June). Locating objects on a latent trait using Rasch analysis of experts' judgments. In *conference "Probabilistic Models for Measurement in Education, Psychology, Social Science and Health", Copenhagen, Denmark]. Retrieved February* (Vol. 20, p. 2026).

Crompvoets, E. A., Béguin, A. A., & Sijtsma, K. (2020). Adaptive pairwise comparison for educational measurement. *Journal of Educational and Behavioral Statistics, 45*(3), 316–338.

Frick, H., Strobl, C., Leisch, F., & Zeileis, A. (2012). Flexible Rasch mixture models with package psychomix. *Journal of Statistical Software, 48*(7), 1–25.

Han, S., Rijmen, F., Boykin, A. A., & Lottridge, S. (2025). Leveraging fine-tuned large language models in item parameter prediction. In *Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con),* Vol. 1: Full Papers (pp. 250–264). National Council on Measurement in Education.

Kolesnikova, D., Fedyanin, K., Hofman, A. D., Brinkhuis, M. J., & Bolsinova, M. (2026). Estimating Item Difficulty with Large Language Models as Experts. *arXiv preprint arXiv:2605.18562*.

Li, M., Jiao, H., Zhou, T., Zhang, N., Peters, S., & Lissitz, R. W. (2025). Item difficulty modeling using fine-tuned small and large language models. In *Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con),* Vol. 3: Coordinated Session Papers (pp. 48–55). National Council on Measurement in Education.

Mislevy, R. J., & Verhelst, N. (1990). Modeling item responses when different subjects employ different solution strategies. *Psychometrika*, *55*(2), 195–215.

Mojoyinola, M., Kehinde, O. J., & Tang, J. (2025). Enhancing item difficulty prediction in large-scale assessment with large language model. In *Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con)*, Vol. 2: Works in Progress (pp. 218–222). National Council on Measurement in Education.

Rost, J. (1990). Rasch models in latent classes: An integration of two approaches to item analysis. *Applied Psychological Measurement, 14*(3), 271–282.

Shi, L., Ma, C., Liang, W., Diao, X., Ma, W., & Vosoughi, S. (2025, December). Judging the judges: A systematic study of position bias in LLM-as-a-judge. In *Proceedings of the 14th International Joint Conference on Natural Language Processing and the 4th Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics* (pp. 292-314).

Zheng, L., Chiang, W.-L., Sheng, Y., Zhuang, S., Wu, Z., Zhuang, Y., Lin, Z., Li, Z., Li, D., Xing, E. P., Zhang, H., Gonzalez, J. E., & Stoica, I. (2023). *Judging LLM-as-a-Judge with MT-Bench and Chatbot Arena*. In *Advances in Neural Information Processing Systems* (Vol. 36). Track on Datasets and Benchmarks.

LLM-As-An-Assessor: Can Open-Weight LLMs Assess Computational Thinking via Student-Designed Embodied Games?

William Lee
University of Massachusetts Amherst
williamlee@cs.umass.edu

Sai Gattupalli
Society & AI Research Group
sai@societyandai.org

Ivon Arroyo
University of Massachusetts Amherst
ivon@cs.umass.edu

Brendan O'Connor
University of Massachusetts Amherst
brenocon@cs.umass.edu

Abstract

We evaluate whether open-weight large language models (LLMs) can assess and infer Computational Thinking (CT) knowledge in middle-school students' finite-state game designs against human-coded labels. Across two game designs, one high-fidelity with rich CT behaviors and one low-fidelity with sparse CT behaviors, models detect some CT behaviors at near-human agreement. However, they also consistently over-credit behaviors that are absent and fail on basic characteristics retrieval, loop/branch detection, and state-level standards mapping. A frontier baseline, Claude Fable 5, matched the human coder perfectly on the rich game but did no better than open-weight models on the sparse one. The largest model is among the least reliable. We trace several failures to fixable task setup and argue for further investigation before such systems inform teacher-facing feedback.

1 Background

Computational Thinking (CT) is now a foundational K-12 competency (Weintrop et al., 2021). Yet teachers lack practical tools to assess it from student work. Using the *WearableLearning* (WL) platform (developed at the Advanced Learning Technologies Lab, UMass Amherst), grades 5-8 students design games as finite state automata (FSA) in a block-based editor in which CT competence is embedded directly in each design's structure (Arroyo et al., 2023). Educators interviewed previously in a pilot study revealed the importance of formative CT assessments aligned to state and national standards, echoing broader evidence that teacher-facing analytics must be grounded in teachers' information needs (Hadi Nezhad et al., 2024). Although LLMs are capable of reading these FSA artifacts, using these models as an assessment instrument demands evidence of validity, reliability, and fairness. We focus on open-weight models because schools need assessment that runs locally, so they can protect student data and eliminate cloud-based cost. However, their assessment abilities for structured CT artifacts remain largely undocumented. Our 2024 pilot (Lee and Arroyo, 2025) showed frontier and reasoning models can match human counts of basic FSA characteristics (F1 $\approx$ 1.0) while advanced analytics (e.g., loop detection) remained fragile.

2 Related Work

Assessment, especially for CT, is an active field including rubrics, artifact analysis, and scoring (Weintrop et al., 2021). Our pilot showed that models can recover basic FSA characteristics (e.g., states) but are sensitive to prompt wording – detecting loops only once the word "loops" was reworded to "cycles" (Lee and Arroyo, 2025). Two literatures inform the present study. First, using LLMs as judges is known to introduce systematic biases and "explanations" could be unfaithful to the model's actual computation (Turpin et al., 2023), which cautions against showing models' rationales to teachers. Second, the operations our task demands have known issues: graph reasoning over structured inputs has been shown to be difficult for LLMs (Wang et al., 2023; Fatemi et al., 2024), and aggressive low-bit quantization, common in local deployment, could further degrade numerical reasoning (Li et al., 2025).

3 Objectives

This research investigates if local open-weight LLMs can assess student-created CT artifacts. **RQ1:** To what extent can open-weight LLMs detect the CT design patterns presented in student FSA games when compared and measured against human-coded labels? **RQ2:** To what extent does CT behavior and standards detection become reliable enough to support valid formative feedback? An additional goal is to identify the evaluation

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 289–297
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

metrics, prompt structures, and rubric designs that make an *LLM-As-An-Assessor* pipeline trustworthy and informative for educators, researchers, and camp counselors.

4 Data

The corpus consists of two student-designed WL games. Each game is represented in a structured JSON format (states, transitions, connections, branching). One game is a high-fidelity design with complex patterns (Fraction Pong) while the other is a low-fidelity design with simple patterns (Square Roots). For each game, a human-coded standard captures 13 CT behavior judgments and the basic characteristics. Basic characteristics are extracted by parsing the JSON via Python. Hence, we found Fraction Pong has 23 states, 36 transitions, 38 connections, 5 cycles, and 7 branches while Square Roots has 12 states, 10 transitions, 11 connections, 0 cycles, and 0 branches. As for gold-standard label reliability, two researchers independently coded the 13 behaviors for two games. Agreement was reached on 84.6% of judgments with Cohen's κ = 0.692 (Cohen, 1960). Square Roots was coded identically (100%) and all disagreement fell on Fraction Pong (69%). Disagreements were resolved by discussion to form the gold labels.

5 Methods

Each model receives an identical payload – the final-design JSON and the WL behavior rubric. Inference runs locally on the open-weight models (`llama.cpp` behind a `llama-swap` proxy; vLLM for larger models) and, for the baseline, the Claude API. We evaluate a set of open-weight LLMs (Table 1) chosen for architectural diversity (dense, sparse mixture-of-experts, and hybrid Mamba–Transformer, about 27B-120B parameters) with a proprietary frontier model, Claude Fable 5, serving as the baseline.

5.1 Pipeline Evolution

Our current evaluation approach has evolved through three phases. In Phase 1, each model received the full game JSON and the rubric together with a single, complex instruction as a prompt through a chat interface (Open WebUI). Results were poor: models frequently regurgitated the prompt back – restating the rubric and instructions rather than analyzing the artifact. The responses produced were fluent and professional but hollow tables. We attribute this likely to overload, since a large structured JSON payload plus a multi-part instruction in one monolithic prompt probably exceeded what most open-weight models could hold and reason over. In Phase 2, following consultation with our advisors, we decomposed the single instruction into a series of seven sub-prompts and added in-context examples to each model. Hence, the model addressed one CT sub-task at a time with a concrete pattern to follow. This improved output quality. In Phase 3, we automated the pipeline through each model's OpenAI-compatible API served through custom Python, `llama.cpp`, and `llama-swap` orchestration, replacing the manual chat interface. This automation enables reproducible runs, captures reasoning traces, and collects responses programmatically. Appendix A shows the Phase 1 prompt in full and summarizes the seven Phase 2 sub-prompts.

5.2 Analysis

Outputs from the LLMs are scored against the human labels. For the 13 presence/absence CT behavior judgments, we use precision, recall, F1 (Manning et al., 2008), and Cohen's κ (Cohen, 1960), which specifically measures agreement beyond what chance would produce. We report per-model metrics, an item-by-model agreement heatmap, and per-model confusion matrices.

6 Results

6.1 CT Behavior Detection

Both games are scored on the same 13 behaviors against gold labels determined by the two independent coders. On the high-fidelity Fraction Pong game, 10 of 12 open-weight models gave usable output (Nemotron-3-nano and Phi-4-reasoning-plus unfortunately did not produce responses consistently). Their F1 ranged from 0.31 to 0.95 (mean $\approx$ 0.80; accuracy $\approx$ 0.74; $\kappa \approx$ 0.45). Gemma-4-31B was the best open-weight model (F1 = 0.95, κ = 0.75), and the frontier baseline matched the gold exactly (κ = 1.00). On the low-fidelity Square Roots game, agreement fell (average F1 $\approx$ 0.59, range 0.50 – 0.67; average $\kappa \approx$ 0.37). Table 2 illustrates the per-model values.

An important finding is about how many behaviors the models say are present. Fraction Pong has 11 of 13 behaviors present and Square Roots has 4. The models did not track this difference. On average these models detected about 7.9 be-

haviors as present on Fraction Pong and about 5.3 on Square Roots. For the rich game, this means these models seem to miss detecting real CT behaviors, whereas on the sparse game, it means these models credit CT behaviors that are not actually there. The confusion matrices show this directly (Figure 2). With Fraction Pong, every model had precision = 1.00: pooled across models there were zero false positives and 34 misses. The pattern flipped for Square Roots: 31 false positives and 15 misses. We think one reason for this difference is the design of the games (Fraction Pong has only two absent CT items, and Square Roots only four present CT items). The errors also cluster on specific CT items (Figure 1). For Fraction Pong, 7 of 11 models missed both "Encouraging feedback" and "Difficulties/levels," while all models got "Sequences" correct. For Square Roots, every model credited "Outputs: text," which the game lacks, and every model missed "Inputs: buttons," which it has. When most models make the same mistake, the cause is possibly in the rubric or the prompt.

6.2 Basic CT Characteristics Retrieval

When counting Fraction Pong's basic characteristics, models did not perform well. The gold has 23 states, 36 transitions, 38 connections, 5 cycles, and 7 branches. Models undercounted these characteristics heavily. The closest, DeepSeek-R1-distill, reported 21 states, 20 transitions, 23 connections, 1 cycle, and 3 branches while several models reported zero transitions or connections. No model found all 5 cycles, which repeats the "loop-detection" failure from our pilot study. Square Roots has no cycles or branches, so it cannot test loop detection.

6.3 State-Level Standards Mapping

Earlier in the evolution of our evaluation pipeline, we attempted to see if these models can extract CT behaviors and map them to the Massachusetts Digital Literacy and Computer Science (DLCS) standards. Models regurgitated the text description of a DLCS standard instead of making any inference. We traced this to the rubric and possibly the prompt: many of our rubric items mixed the DLCS standard codes and the descriptions in one field and stacked several standards behind line breaks when feeding into the prompt. Two of eight models on Fraction Pong returned output we could not parse (for example, raw Python). The polished look of many generated Markdown tables hid these errors.

6.4 Frontier Baseline

The effectiveness of the frontier model used in this evaluation depended on the richness of a game's design. Claude Fable 5 matched the human gold labels exactly (F1 = 1.00, κ = 1.00) for Fraction Pong whereas the best open-weight model only reached κ = 0.75. For Square Roots, Fable scored F1 = 0.60 and κ = 0.37 which is in the middle of the pack, below the three best open-weight models (RNJ-1, Granite-4-tiny, and Qwen3.6, each κ = 0.49). Interestingly, Fable's errors on Square Roots were the same ones every other model made: it credited "Outputs: text" and "Inputs: text," which the game did not have, and missed "Inputs: buttons."

Hence, it seems that the frontier model was able to find every behavior that was present in a richly designed game. Yet, the frontier model, just like the other models, has difficulty confirming behaviors that were *absent* in a sparsely designed game. Rankings among open-weight models also did not carry over between games. RNJ-1 was the best open-weight model on Square Roots (κ = 0.49) and the worst on Fraction Pong (κ = 0.06, missing 9 of 11 present behaviors).

The largest model, GPT-OSS-120B, was the least reliable on Square Roots (κ = 0.25) and only performed decently on Fraction Pong (κ = 0.35). Neither being a frontier model nor having a large number of parameters predicts reliable assessment across games.

7 Lessons Learned

Beyond the quantitative agreement metrics, systematic inspection of model outputs revealed some interesting patterns. We group them into five categories. (1) Prompt regurgitation: many models often repeat the rubric or instruction instead of analyzing the artifact. This was especially common during Phase 1's monolithic setup. (2) Description vs. standard mapping: models return a DLCS standard's prose description rather than mapping a standard to a CT behavior, which also happened in Phase 1. (3) Count contradictions: counts of basic characteristics disagree with the game's JSON structures suggesting that these models might have difficulties parsing messy and complex JSON objects. (4) Unparseable output: raw HTML, Python code, or malformed text was produced in place of the explicitly requested judgment, mapping, or inference (two of eight models on Fraction Pong). (5) Non-response: no usable output at all was returned

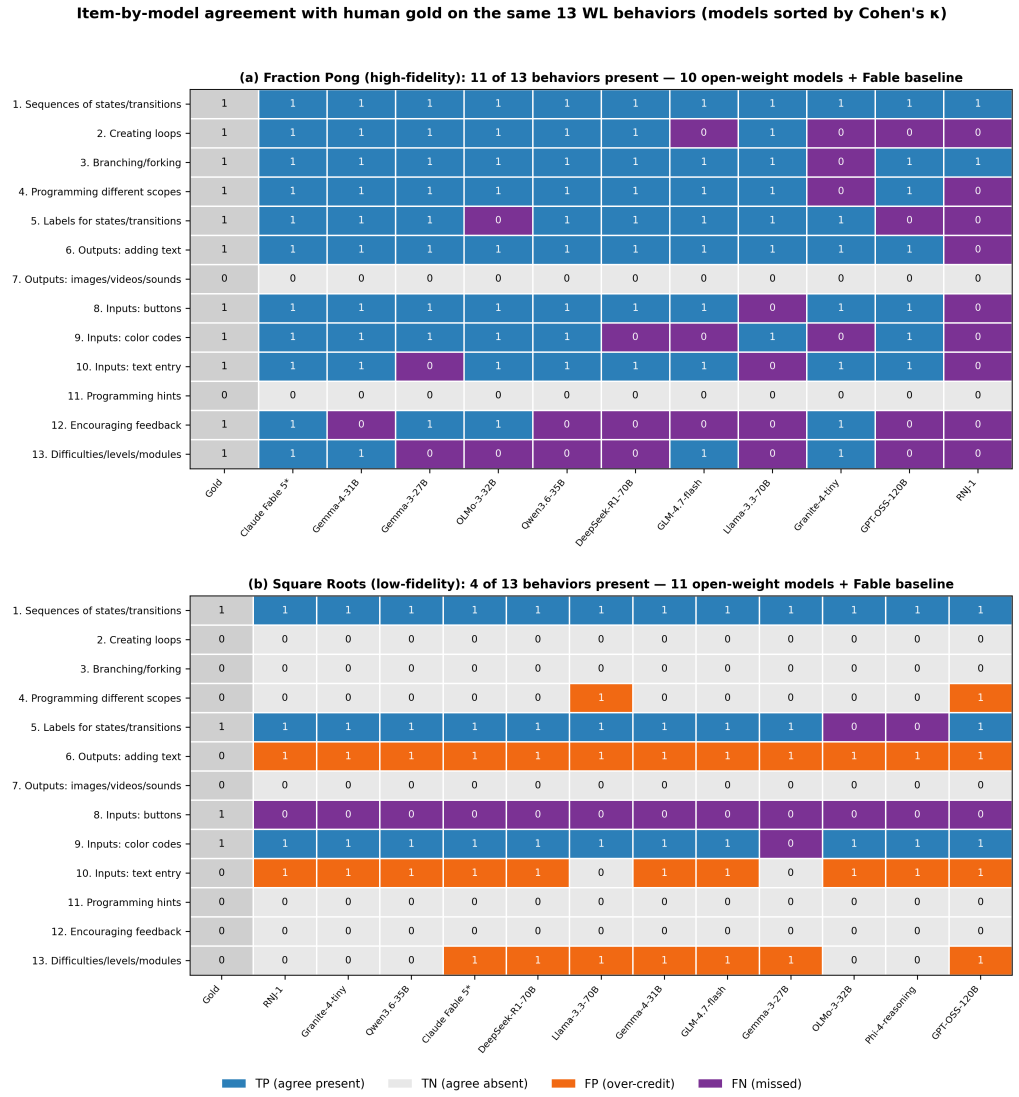


Figure 1: Item-by-model agreement with human gold on the same 13 WL behaviors, models sorted by κ. Top: Fraction Pong (high-fidelity; 11 of 13 behaviors present; 10 open-weight models + Fable baseline). Bottom: Square Roots (low-fidelity; 4 of 13 present; 11 open-weight models + Fable baseline). Orange = over-crediting (FP); purple = missed (FN).

(e.g., Nemotron-3-nano on Square Roots).

A common observation ties these to our instrumentation premise. Very early on, nearly all outputs were rendered as well-formed Markdown tables as requested. They looked authoritative, with consistent columns, plausible language, and a confident tone. Yet, when inspected carefully, the underlying content was frequently wrong, hollow, or just regurgitation from the prompt. We believe a well-structured table of wrong answers would be uninformative, to say the least, for a teacher compared with a broken one because it invites trust. This "rendering–substance" gap was most severe in Phase 1, when a full JSON payload and a single monolithic prompt overwhelmed most of the models. Reducing the prompt through decomposition and in-context examples led to the quality gains we saw in the later evolution of the pipeline.

8 Implications, Discussion, and Conclusions

Several implications follow. First, several failures are setup- and configuration-related and not likely related to the models' ability. For example, a rubric schema we created that fuses DLCS codes with descriptions, and a bundled prompt without any in-context examples likely led to task performance degradation and interference. Each is addressable through schema cleanup, task decomposition, and in-context examples. The revised pipeline is being

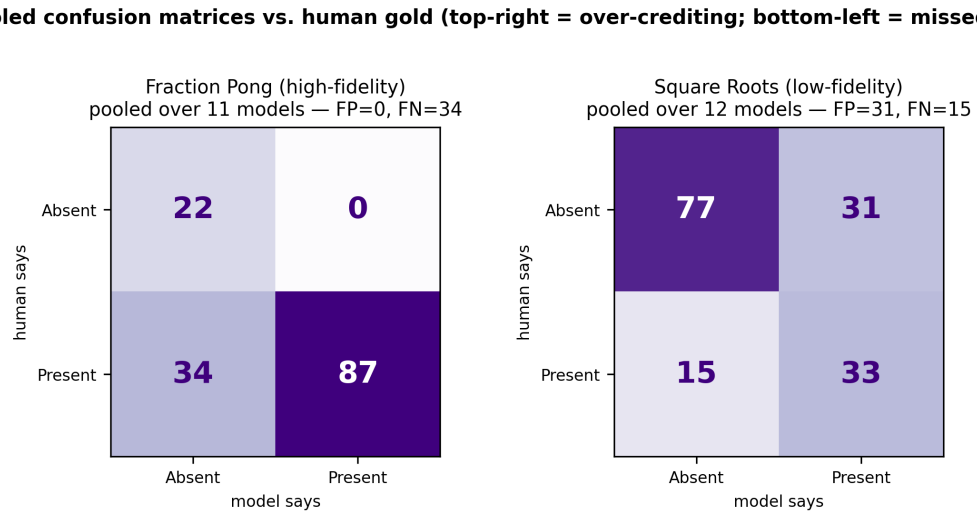


Figure 2: Pooled confusion matrices against human gold. Fraction Pong shows only misses (0 FP, 34 FN); Square Roots is dominated by over-crediting (31 FP, 15 FN). Top-right = over-crediting; bottom-left = missed. Pooled counts are descriptive only.

re-run with additional game designs.

Second, the over-crediting pattern is a problematic teacher-facing risk. Under a sparse design, some of our models credit behaviors the student did not demonstrate. So a "what's missing" feedback, which depends on correctly detecting absence, would tell a teacher a student did things they actually did not do. On the other hand, for a rich design, the model misses behaviors the student did show, so it fails to recognize real skill and learning. Because detection reliability varies by behavior (Figure 1) and by design fidelity, we think absence-based feedback should be surfaced only for behaviors a model detects reliably. The absence problem fits what is known about LLMs unfortunately. These models tend to affirm what is in front of them and handle negation and absence poorly (Truong et al., 2023). That is why even the frontier model, which found every present behavior in the rich game, credited absent ones in the sparse game. Therefore, we believe the frontier result does not indicate whether open or closed models are better. Instead, it shows that the advantage depends on the design and the raw content. Open-weight models that run locally could still be a good fit for schools once the whole assessment setup, not just the model, has been validated.

Third, our results raise the question of whether different levels of game design call for different models, and whether teachers should be able to choose which model assesses which design. Our data caution against a teacher-facing model picker. Rankings did not transfer between the two games, so a choice made on one design would not hold on another, and with only two games, the rankings themselves rest on a small number of samples. A model picker would also move a technical burden and decision onto teachers and could invite choosing whichever model gives the preferred answer. Two future alternatives might worth considering. One is to run several models on a game design and treat their agreement as a confidence signal: behaviors that most models agree on are surfaced, and disagreements are flagged for the teacher. The other is to adjust for the design's density automatically. A design's basic characteristics (states, transitions, connections) can be counted directly from the JSON before any model runs, so instead of asking a teacher to guess how complex a design is and choose a model to match, let the pipeline measure the complexity itself from the JSON and adjust its own skepticism accordingly.

Fourth, everything reported here assesses a final product. We attempted to assess progression across the two-day WL curriculum but found that the two-day log data need validation before they can serve as indicators of progression. Monitoring progress in a classroom is a natural goal for an LLM-based tool, but our results suggest something else. Progress is a difference between snapshots, so it inherits the error of each snapshots. Yet the most direct measures of progress, changes in states, transitions, connections, and cycles, might not need a model at all. They can be computed determin-

istically from each saved JSON mentional previously. A future design would then track structural progress by parsing for these basic characteristics then apply the LLM for interpreting those changes in CT terms and providing teachers with feedback.

Our results so far cover two games with two coders whose disagreements were resolved by discussion. Hence, the number of game designs and CT behaviors is relatively small, so per-model rankings and confusion matrices should be considered as descriptive statistics. Rankings among models are also unstable. They do not carry over between the two games. For example, the RNJ-1 model ranked first on one game design and last on the other. Loop and branch detection can only be tested on the high-fidelity game. Finally, because our two games sit at opposite ends, one very rich and one very sparse, we also cannot generalize our findings yet. More game designs with a balanced number of present and absent behaviors are planned to be tested in the future.

In summary, we think open-weight LLMs can extract and infer some CT behaviors at near-human agreement but are not yet reliable for the counting, loop/branch detection, and standards mapping that trustworthy formative assessment requires. Especially for simple designs, they seem to over-credit and sometimes fabricate evidence. And the tidy format of their output and tables hides errors and regurgitation. Our contributions are (1) a measurement-grounded test harness (F1 and κ for behavior detection, error and element-level analysis for basic characteristics); (2) practical prompt- and rubric-engineering guidance; and (3) evidence, including a frontier baseline, that locally deployable open models can be competitive. This research essentially illustrates asking the question "under what conditions are these open-weight LLMs good enough for assessment?" is better than merely "are open models capable?"

Model	Params	Architecture	Reasoning
DeepSeek-R1-distill-70B	70B	Dense (Llama-3.3 base, R1 distill.)	Yes
Llama-3.3-70B-Instruct	70B	Dense	No
Gemma-3-27B	27B	Dense (vision-capable)	No
Gemma-4-31B	∼31B	Dense	No
RNJ-1 (Essential AI)	—	—	—
Nemotron-3-nano (NVIDIA)	Nano	Hybrid Mamba–Transformer	Yes
OLMo-3-32B-think (AllenAI)	32B	Dense	Yes
Phi-4-reasoning-plus	∼14B	Dense	Yes
Granite-4-h-tiny (IBM)	Tiny	Hybrid Mamba–Transformer	No
Qwen3.6-35B-A3B	35B / ∼3B	Sparse MoE	Yes
GLM-4.7-flash	—	MoE —	—
GPT-OSS-120B (OpenAI)	∼120B / ∼5B	Sparse MoE	Yes
Claude Fable 5 (baseline)	—	Proprietary	Yes

Table 1: Model roster. Some specifications for recent releases require verification against official model cards.

	Fraction Pong (high-fid.)		Square Roots (low-fid.)	
Model	F1	κ	F1	κ
Gemma-4-31B	0.95	0.75	0.60	0.37
Gemma-3-27B	0.90	0.58	0.50	0.28
OLMo-3-32B	0.90	0.58	0.50	0.28
Qwen3.6-35B	0.90	0.58	0.67	0.49
DeepSeek-R1-70B	0.84	0.45	0.60	0.37
GLM-4.7-flash	0.84	0.45	0.60	0.37
Llama-3.3-70B	0.78	0.35	0.60	0.37
Granite-4-tiny	0.78	0.35	0.67	0.49
GPT-OSS-120B	0.78	0.35	0.55	0.25
RNJ-1	0.31	0.06	0.67	0.49
Nemotron-3-nano	n/r	n/r	n/r	n/r
Phi-4-reasoning	n/r	n/r	0.50	0.28
Claude Fable 5 (baseline)	1.00	1.00	0.60	0.37
Mean	**0.80**	**0.45**	**0.59**	**0.37**

Table 2: Behavior-detection agreement with human gold (F1 and Cohen's κ) on both games, scored on the same 13 behaviors against two-coder resolved gold. Open-weight rows sorted by Fraction Pong κ. n/r = non-response.

#	Task	Behaviors covered	Words	Worked example
1	Extract and count basic characteristics from the JSON only	states, transitions, connections, cycles, branches	830	Yes (full extraction output of another game)
2	Mark presence/absence of structural behaviors	sequences; loops; branching/forking	197	Yes
3	Scopes and labels	programming different scopes; labels for states/transitions	208	Yes (state ID + description)
4	Output behaviors	outputs: adding text; outputs: images/videos/sounds	257	Yes (text and media examples)
5	Input behaviors	inputs: buttons; color codes; text entry	297	Yes (transition IDs)
6	Support behaviors	programming hints; encouraging feedback	317	Yes (from other games)
7	Systems-level organization	difficulties/levels/modules	208	Yes (Level 1–3 example)

Table 3: The seven Phase 2 sub-prompts. Each targets a few behaviors and includes a worked example; all but the first also ask for a rendered presence/absence table with justifications from the JSON.

References

Ivon Arroyo, Avery H. Closser, Francisco Castro, Hannah Smith, Erin Ottmar, and Matthew Micciolo. 2023. The WearableLearning platform: A computational thinking tool supporting game design and active play. *Technology, Knowledge and Learning*, 28(4):1815–1824.

Jacob Cohen. 1960. A coefficient of agreement for nominal scales. *Educational and Psychological Measurement*, 20(1):37–46.

Bahare Fatemi, Jonathan Halcrow, and Bryan Perozzi. 2024. Talk like a graph: Encoding graphs for large language models. In *Proceedings of the 12th International Conference on Learning Representations (ICLR)*.

Mohammad Hadi Nezhad, Francisco Castro, Beverly Woolf, and Ivon Arroyo. 2024. Math teachers' in-class information needs and usage for effective design of classroom orchestration tools. In *Technology Enhanced Learning for Inclusive and Equitable Quality Education (EC-TEL 2024)*, volume 15159 of *Lecture Notes in Computer Science*, pages 299–314. Springer.

William Lee and Ivon Arroyo. 2025. Toward extracting computational thinking evidence with large language models: Empowering K–12 educators. In *Artificial Intelligence in Education (AIED 2025)*, volume 2590 of *Communications in Computer and Information Science*. Springer.

Zhen Li, Yupeng Su, Runming Yang, Congkai Xie, Zheng Wang, Zhongwei Xie, Ngai Wong, and Hongxia Yang. 2025. Quantization meets reasoning: Exploring LLM low-bit quantization degradation for mathematical reasoning. arXiv:2501.03035. *Preprint*, arXiv:2501.03035.

Christopher D. Manning, Prabhakar Raghavan, and Hinrich Schütze. 2008. *Introduction to Information Retrieval*. Cambridge University Press.

Thinh Hung Truong, Timothy Baldwin, Karin Verspoor, and Trevor Cohn. 2023. Language models are not naysayers: An analysis of language models on negation benchmarks. In *Proceedings of the 12th Joint Conference on Lexical and Computational Semantics (*SEM 2023)*, pages 101–114. Association for Computational Linguistics.

Miles Turpin, Julian Michael, Ethan Perez, and Samuel R. Bowman. 2023. Language models don't always say what they think: Unfaithful explanations in chain-of-thought prompting. In *Advances in Neural Information Processing Systems*, volume 36.

Heng Wang, Shangbin Feng, Tianxing He, Zhaoxuan Tan, Xiaochuang Han, and Yulia Tsvetkov. 2023. Can language models solve graph problems in natural language? In *Advances in Neural Information Processing Systems*, volume 36.

David Weintrop, Daisy Wise Rutstein, Marie Bienkowski, and Steven McGee. 2021. Assessing computational thinking: An overview of the field. *Computer Science Education*, 31(2):113–116.

A Prompts

This appendix shows the single Phase 1 prompt in full and summarizes the seven Phase 2 sub-prompts (also used, automated, in Phase 3). The contrast illustrates the change in task design: one short instruction pointing at an entire rubric spreadsheet, the game JSON, and a two-day log, versus seven focused prompts of 200–830 words, each targeting a few behaviors and each containing a worked example drawn from a real game (Table 3). Student-identifying content has been removed.

A.1 Phase 1: single monolithic prompt (141 words)

```
You are a 6th grade computer science
teacher in Massachusetts. Your job is
to use the "CT breakdown in WL
-12_02_2025.xlsx" as a rubric to assess
the computational thinking knowledge
and skills based on a student designed
CT game. Column A of the spreadsheet
represents various behaviors during the
design of a CT game. Columns E and G
are Massachusetts CT standards mapped
to those behaviors.

The CT game is designed as a finite
state automata. The
Fraction_Pong_final.json is the
finalized version of the CT game. The
Fraction_Pong_2_days.csv is a breakdown
of the progress of the design of the CT
game during a 2-day period.

Using the rubric, assess the CT game
complexity with a score between 1 to 5
and provide explanation. Also assess
the student progression using the
Fraction_Pong_2_days.csv. Identify if
the student was making progress and
provide rationale.
```

A.2 Phase 2: seven decomposed sub-prompts

As an example of the decomposed style, sub-prompt 2 reads as follows (example and the rendering instruction included).

```
Second, with the basic characteristics
of the game's design identified and
counted, your next task is to identify
the presence or absence of these basic
characteristics as behaviors. For
example, if there are any occurrences
of states, connections, and transitions
identified in the design, a behavior
called "Creating sequences of states
and transitions" would be marked as 1
```

to indicate the presence of this behavior. Otherwise mark as empty to indicate absence. If there are any occurrences of cycles or loops identified in the design, a behavior called “Creating Loops” would be marked as 1 to indicate the presence of this second behavior. Otherwise mark as empty to indicate absence. If there are any occurrences of branching/forking identified in the design, a behavior called “Branching/Forking” would be marked as 1 to indicate the presence of this third behavior. Otherwise mark as empty to indicate absence.

Also, create and render a table in markdown or HTML format to show the presence or absence of these three behaviors. I must be able to preview the table in a web browser and I do not want it in any raw code. In the table, provide examples and justifications from the “Fraction_Pong_final.json”.

Funnel Plot Analysis of Examinee Performance in Longitudinal Assessment

Aquia Richburg[1], Marcus Walker[2]
[1]aquiar@nccpa.net, [2]marcusw@nccpa.net
National Commission on Certification of Physician Assistants (NCCPA)

Abstract

We identify content areas where examinees underperform in a longitudinal assessment for a high-stakes recertification exam. Using modified funnel plots with a moving baseline to account for item difficulty, we rank topics by relative performance. Preliminary results suggest reasons beyond item difficulty, informing future analyses supporting potential educational interventions.

1 Introduction

Research has shown that self-testing one's own knowledge improves long-term retention over simply restudying material in what is known as the testing effect (Roediger and Karpicke, 2006). Reflecting this principle, healthcare certification programs have increasingly adopted longitudinal assessment (LA) as an alternative to single-point-in-time examinations, with studies suggesting that LA supports continuous learning, identifies knowledge gaps, and improves performance while maintaining assessment of competence (Fraundorf et al., 2023; Giron et al., 2021; Ward et al., 2023). As LA becomes more widely implemented, understanding patterns of examinee performance across content domains may provide valuable insights for optimizing lifelong learning and informing assessment design.

2 Background

Analyzing the change in student performance over time is an important component in measurement theory. Understanding the trajectories of different subpopulations provides insights informing decisions on pedagogical and testing designs. One application of this analysis is the identification of knowledge gaps. These can be thought of as the discrepancy of an individual's ability in relation to an expected performance threshold and the reasons for these gaps can come in many forms. Schmidt (2020) outlines an extensive taxonomy of knowledge gaps and potential strategies to address them. This taxonomy ranges from more surface level, such as understanding textual content, to more abstract, such as requiring implicit context to arrive at a conclusion. In the field of medical licensure, the more metacognitive categories of knowledge gaps could be a better theory for explaining performance difficulties due to the multilayer complexity of medical-domain exam items.

Research in measuring growth trajectories has taken a variety of different forms. Early works like Muthen and Khoo (1998) generalize random coefficient models to incorporate latent variables that can more directly model individual variation in growth parameters between multiple subpopulations. Latent variable frameworks have continued to be popular in recent research where various methods have been proposed to measure and predict population-dependent growth trajectories and in some cases to identify points for intervention (von Davier et al., 2011; Kwon and Zhang, 2025; Liu and Culpepper, 2025; Yamaguchi et al., 2026). While these prior studies primarily use repeated measurements of performance over a longitudinal period, our work differs by also incorporating longitudinal performance of items administered to address deficits in knowledge.

Natural language processing (NLP) has become a powerful tool in the biomedical domain. To address the distributional shift in text from a general domain, several language models (Alsentzer et al., 2019; Lee et al., 2020; Rasmy et al., 2021) have been fine-tuned for tasks such as information extraction from domain-specific or structured text. In more recent work, large language models (LLMs) have been integrated into systems to reformulate unstructured biomedical

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 298–305
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

textual data into formats more readily usable (Wiest et al., 2025; Hsu and Roberts, 2025). Item difficulty prediction has also become a task that increasingly relies on NLP techniques. Contreras (2018) present methods to measure the relative frequency in terms between topic-specific and general texts in an attempt to quantify differences in mental workload among examinees. Ha et al. (2019) construct a set of linguistic and information retrieval features for item difficulty prediction using a corpus constructed from the United States Medical Licensing Examination® (USMLE®). The BEA (Building Educational Applications) 2024 shared task (Shardlow et al., 2024) focuses on automatic difficulty prediction and response time prediction using the USMLE® where the best strategy relied on transformer-based text embeddings. Li et al. (2025) achieved a slightly better model than the best shared task participant by incorporating data augmentation alongside language models.

3 Aims

The main goal of this study is identifying global patterns on a longitudinal assessment using the response data over a cohort of examinees. The exam blueprint consists of several main content areas that are comprised of various topics. During each window of administration, an examinee is presented with a set of items representatively sampled from the distribution of topics. Over the course of the examination, the set of items changes according to response behavior, favoring more items from topics an examinee missed. Identifying topic areas of the blueprint where examinees are more likely to have difficulty than expected could better support the goal of continuous learning.

4 Data

We used the response data for a longitudinal exam program consisting of 12 administration windows (quarters) from 2023 to 2025. Each quarter examinees are provided with 25 items across 14 content areas that can target one of four cognitive tasks. Examinee responses to items are categorized as an original attempt, items drawn from a pool according to blueprint specifications, or a targeted attempt, items chosen according to blueprint specifications **and** addressing a topic encountered in a prior quarter according to an index that incorporates whether an item was answered incorrectly, the confidence of an examinee on that topic and the examinee's perceived relevance of that topic to their clinical practice. Higher index scores increase the likelihood that the topic will appear in future quarters. Each examinee is tested over 12 quarters and may exit early by passing 8 quarters. We hypothesize that low performance on targeted attempts relative to original attempts could signify gaps in an examinee's knowledge on specific topics.

5 Methods

5.1 Confidence and Targeted rankings

The structure of this longitudinal exam is semi-adaptive; the items administered to examinees is informed by information from prior windows of information. After the completion of each item, an examinee is asked to rate their confidence in answering the item and rate the perceived relevance of the item to their practice where higher values indicate stronger agreement. In the attempt to serve the most useful items to examinees, a targeted index is computed incorporating confidence, relevance and the binary performance on the item. The exam prioritizes selecting subsequent items from topics where examinees incorrectly answer the item and have high targeted index.

Confidence ratings and targeted index scores provide direct signals from examinees interacting with the exam; thus ranking items along these dimensions could indicate areas where examinees struggle. For each topic, we compute the mean confidence rating and mean targeted index across all responses from 2023 to 2025. We rank confidence from lowest to highest and rank targeted index from highest to lowest. These rankings are included as comparisons against our funnel ranking which we will describe in the next section.

5.2 Funnel ranking

The main methodology we incorporate is the statistical funnel plot. A funnel plot allows for the comparison of different indicators where the precision of the indicator varies. Spiegelhalter (Spiegelhalter, 2005; Spiegelhalter et al., 2012) utilized funnel plots to propose an alternative to rank medical institutions. Instead of comparing performance indicators using a common confidence interval, they accommodate differences

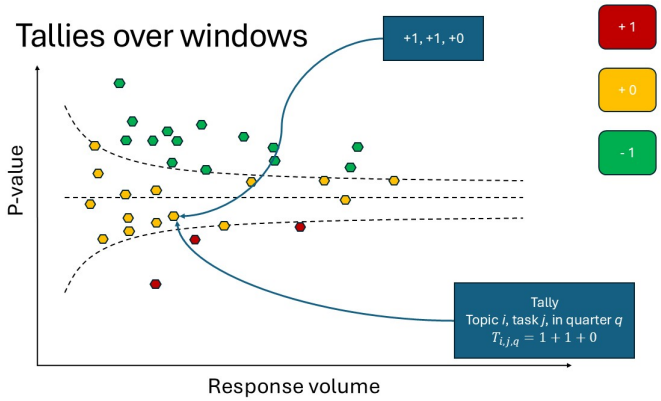


Figure 1: Graphic illustrating the tallying of scores across funnel plot windows. This example demonstrates the tally accrued over three downward consecutive funnel plot windows in funnel window $n = 3$.

in precision by varying the confidence interval to allow for differences in indicator variance. In turn, this creates a funnel-like shape around a chosen baseline (which could be the mean of the indicator) as the confidence interval shrinks, where precision increases. Points that fall within the control limits are not statistically different from each other while those outside have reason for scrutiny.

Let $a_{i,j,q,h}$ be an item attempt that consists of a topic i and task j that is administered in quarter q to examinee h. We define the P-value of a topic-task-quarter combination, $p_{i,j,q}$, as the ratio of all correct responses to incorrect responses for attempts $a_{t,j,q,h}$ across all examinees. We assign a mean difficulty, $\mu_{b_{t,j,q}}$, for the combination of topic-task-quarter, $p_{i,j,q}$, by computing the mean item difficulty parameter $\mathbf{b}$ from the Rasch model across the items of topic i and task j in quarter q. Thus, our points of interest are topic-task combinations across quarters, and the indicator is the associated P-value. Since performance on items is dependent on its difficulty, a common baseline P-value over all items would disadvantage items towards the tail ends of the distribution, since they are further from the mean.

Let A_O and A_T be the collection of $p_{i,j,q}$ computed from original attempts and targeted attempts, respectively. Using the distribution of mean difficulties, $\mu_{b_{i,j,q}}$, we first partition the collection of all type-task-quarter combinations ($p_{i,j,q} \in A_O \cup A_T$) using k-means. This generates N classes of $p_{i,j,q} \in A_O$ for comparisons against $p_{i,j,q} \in A_T$. For our study, we set $N = 5$. Within the structure of the longitudinal exam, examinees usually enter in the first quarter of a calendar year generating a cohort that is administered items from a common pool. In turn, targeted attempts are pulled from the same common pool for subsequent quarters. For example, the 2023 cohort will see original attempt items from the 2023 pool in quarters 1 through 4; then the same cohort will see targeted attempt items from the same pool in quarters 2 through 5 (quarter 5 is equivalent to quarter 1 of 2024). To keep comparisons of items consistent within a pool we partition both A_O and A_T into years $y = 1, 2, 3$ for 2023, 2024 and 2025, respectively. Since targeted attempts are administered after original attempts, the years are staggered by a quarter, i.e. $y = 1$ consists for quarters 1-4 for A_O and 2-5 for A_T. For each partition, we define a baseline P-value as the median of the mean difficulties, $\mu_{b_{i,j,q}}$, from $p_{i,j,q} \in A_O^y$. We choose the median value to avoid possibly skewed distributions of mean difficulties. The baseline P-values are then used as points of comparison against $p_{i,j,q}$ from the corresponding A_T^y.

We plot $p_{i,j,q} \in A_T^y$ for each $y \in (1, 2, 3)$ where the volume of responses and the P-value represent the x- and y-axis, respectively. Using the baselines defined above, we plot $n = N$ funnel plots overlapping the data. For each type-task-quarter combination, we tally scores, $T_{i,j,q,n}$, by assigning -1, 0 or 1 according to whether the point $p_{i,j,q}$ lies above, within or below the control limits for each funnel plot, respectively. We then sum the scores $T_{i,j,q,n}$ across all N funnel plots for each quarter q. This generates a score for each topic-task-quarter combination. See figure 1 for a graphic depicting the process.

We compute final scores, $S_{i,j}$, for topic-task combinations as the mean $T_{i,j,q}$ over all quarters;

and for topics, S_i, as the mean score $S_{i,j}$ over tasks in topic i. Ultimately, we rank all S_i or $S_{i,j}$ from highest to lowest where higher scores indicate areas where examinees tend to have more difficulty. We allow ties to prevent spurious rankings. The sequence of ranking equations are listed below:

$$T_{i,j,q} = \sum_{n=1}^{N} T_{i,j,q,n}, \quad (1)$$

$$S_{i,j} = \frac{1}{Q} \sum_{q=1}^{Q} T_{i,j,q}, \quad (2)$$

$$S_i = \frac{1}{m_j} \sum_{m=1}^{m_j} S_{i,j} \quad (3)$$

where N is the number of funnel plot windows, Q is the total number of quarters and $m_j \in [1, 4]$ is the number of tasks in topic i.

6 Results

Within each of the 14 content categories, we list the top ranked topics in Table 1 along with its allocation according to the blueprint[1]. Unsurprisingly, the highest-ranked topic is from the cardiovascular system, given its allocation is the highest at 12%. The remaining ranks do not appear to follow a pattern; the genitourinary system is only 5% but has a topic at rank 3 while gastrointestinal and pulmonary (both 10%) have ranks at 13 and 18, respectively. Similarly, item difficulty does not have a clear connection to rank; compare rank 1 with difficulty 1.262 and rank 18 with difficulty 2.038. We also include the overall top 10 funnel-ranked topics in Table 3 and the highest- and lowest-ranked topics for each ranking type in Table 2.

In figure 2, we plot the correlation of funnel, confidence and index rankings against mean topic difficulty. All three plots show a negative correlation against item difficulty; however confidence and index rankings have tighter groupings of items across rank levels compared to the funnel rankings. Quantitatively, we compute the Kendall τ correlation as -0.416, -0.566 and -0.705 for funnel, confidence and index rankings, respectively. This indicates that the funnel rankings describe trends in examinee behavior that are not explained by item difficulty.

For the topic with the highest mean difficulty (hypertension, secondary in Cardiovascular System with difficulty 3.527), there is a large difference between the funnel rank at 80 and confidence and index ranks at 15 and 17, respectively. In contrast the lowest mean difficulty topic (rabies in Infectious Diseases with difficulty -3.238) has a slightly smaller difference between ranks with 181, 237 and 237 for funnel, confidence and index ranks, respectively. The greatest difference in funnel and index ranks is $225 - 82 = 143$ for pneumonia: bacterial, fungal, viral (Pulmonary System) and the greatest difference in funnel and confidence ranks is $212 - 31 = 181$ for valvular disorders: regurgitation, prolapse, stenosis (Cardiovascular System). We also compute Kendall τ as 0.381 for funnel against confidence ranks, 0.497 for funnel against index ranks and 0.608 for confidence against index ranks.

7 Discussion and Conclusion

Our main goal for this study was to use the response data from a high-stakes longitudinal recertification exam to identify topics where examinees particularly struggle. Within the structure of the exam, there are survey questions that ask about each item's relevance to the examinee and the confidence in their answer. Although these self-reported scores provide a signal for identifying topics that might be more difficult than others, there is a risk of bias that might distort the actual difficulty seen in the data. By utilizing the more objective response data, we can more directly measure the variation in performance across different areas of the blueprint.

Our results show there is little consensus in relative rankings between the three ranking strategies which is somewhat surprising in the case of confidence versus index rankings given that examinee confidence is incorporated in the calculation of the targeted index. Pairwise correlations are not strong especially between funnel rankings and the other two. The median difference between funnel and confidence rankings and funnel and index rankings is -1 and 3, respectively. This means that topics ranked by index tend to be higher than when ranked by funnel whereas funnel and confidence rankings are

[1]We exclude Emergent Topics (Legal, Ethical, DEI) at 2%.

Medical content category	Allocation	Topic	Funnel rank	Mean difficulty
Cardiovascular System	12%	Aortic aneurysm, abdominal	1	1.262
Dermatologic System	5%	Erythema multiforme	24	1.327
Endocrine System	8%	Type 2 diabetes mellitus/hyperosmolar hyperglycemic syndrome	4	2.765
Eyes, Ears, Nose & Throat	8%	Eustachian tube dysfunction	7	0.982
Gastrointestinal System/Nutrition	10%	Constipation	13	1.695
Genitourinary System	5%	Benign prostatic hyperplasia	3	0.895
Hematologic System	4%	Hypocoagulable states	5	0.495
Infectious Diseases	7%	Septic disorders: sepsis, septic shock, systemic inflammatory response syndrome	13	1.871
Musculoskeletal System	8%	Osteoarthritis	13	1.612
Neurologic System	5%	Major neurocognitive disorder/dementia	7	0.732
Psychiatry/Behavioral Science	7%	Disruptive, impulse-control, and conduct disorders: oppositional defiant disorder, conduct disorder	20	1.379
Pulmonary System	10%	Asthma	18	2.038
Renal System	4%	Renal cell carcinoma	19	1.558
Reproductive System	5%	Pregnancy loss	16	0.594

Table 1: List of the top ranked topics per organ system including the mean topic difficulty and each organ system's percent allocation on the blueprint. We highlight the easiest and most difficult topic by mean difficulty; and the highest rank and lowest rank per topic.

Medical content category	Topic	Funnel rank	Index rank	Confidence rank	Mean difficulty
Cardiovascular System	Aortic aneurysm, abdominal	1	28	82	1.262
Cardiovascular System	Venous insufficiency	1	25	24	1.450
Infectious Diseases	Rubella	34	1	41	2.991
Endocrine System	Thyroiditis	34	19	1	2.082
Neurologic System	Bell palsy	233	120	114	-0.122
Infectious Diseases	Rabies	259	237	237	-3.238

Table 2: The top portion lists the top ranked by each ranking system and the bottom portion the lowest ranked. Note that rabies is the lowest ranked for both confidence and index rankings.

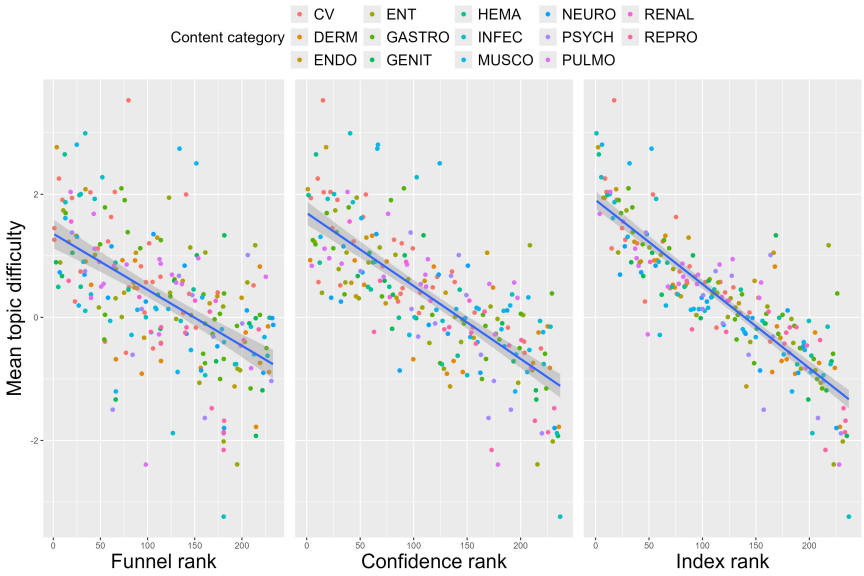


Figure 2: A plot of the correlation of funnel rankings, confidence rankings and index rankings against mean topic difficulty. We also include regression lines over all topics including confidence intervals.

closer on average. Combined with the weak correlation, the different ranking types have some similarity in ordering though there are some topics with large differences in ranking. A natural question arises about how the combined information from the rankings informs us about the behavior of examinees. For example, for a topic that is ranked low by confidence and high by funnel might indicate an area where examinees tend to be overconfident or a similar ranking between confidence and index could indicate self-reported confidence better aligned with performance.

An implication related to this is uncovering different subpopulations within the examinees by their performance in relation to the topic rankings. Currently we do not distinguish responses between different groups of examinees. Since all examinees are given targeted items even if their prior quarter was perfect, this creates the concern that performance on targeted items may be inflated due to the contribution of responses of high performing examinees. Topics with similar rankings across the ranking types might indicate more true knowledge gaps relative to the examinee population. On the other hand, topics with more discordant rankings might suggest weaker or stronger knowledge gaps among different subgroups. One possibility would be to create new rankings according to these separate groups which could be useful to have a more accurate understanding of how groups of examinees are performing. In an operational application, this could lead to different ways of prioritizing topics for emphasis across examinees.

Since each type of ranking has modest at best correlation with item difficulty, further analysis of the items themselves could shed light on the rationale. Traditional NLP methods like feature engineering on textual attributes have shown middling success at categorizing item difficulty (Contreras, 2018; Ha et al., 2019) and even language models fine-tuned for the biomedical domain have difficulty in improving on difficulty prediction compared to general domain models (Li et al., 2025). This suggests there may be a more metacognitive rationale explaining these performance gaps that are not captured in the text itself. Additionally we only analyze rankings at the topic level, separating into topic-task combinations might provide more information that is lost averaging over topics.

In summary, the results of this study provide a new technique for relative comparison of performance across topics within a longitudinal examination. Practical applications such as incorporation into recommender systems or personalized study plans over emphasized topics could lead to overall increased performance or earlier exits from the LA.

Limitations

A major limitation to the current study is the structure of the data. Although there is a healthy

Medical content category	Topic	Funnel rank	Index rank	Confidence rank	Mean difficulty
Cardiovascular System	Venous insufficiency	1	25	24	1.450
Cardiovascular System	Aortic aneurysm, abdominal	1	28	82	1.262
Genitourinary System	Benign prostatic hyperplasia	3	33	62	0.895
Endocrine System	Type 2 diabetes mellitus/hyperosmolar hyperglycemic syndrome	4	2	18	2.765
Hematologic System	Hypocoagulable states	5	71	57	0.495
Cardiovascular System	Coronary artery disease	6	9	10	2.256
Neurologic System	Major neurocognitive disorder/dementia	7	43	52	0.732
Eyes, Ears, Nose & Throat	Eustachian tube dysfunction	7	41	58	0.892
Genitourinary System	Overactive bladder	9	45	48	0.672
Cardiovascular System	Myocardial infarction, acute: non-ST-segment elevation myocardial infarction (NSTEMI)	10	20	29	1.910

Table 3: Top 10 topics overall by funnel rank.

volume of responses across examinees, the longitudinal range is only three years. This limits our comparisons across administration windows and our current trends might not be stable across a longer timescale.

Another limitation is that confidence and relevance scores are self-reported. Since examinees interact with the scale individually reported scores are not necessarily equivalent across examinees. This has the possible impact of skewing the interpretation of the scores where, for example, a topic is judged to have higher average examinee confidence while having relatively low performance. For reliability it would be helpful to equate them on the same scale.

References

Emily Alsentzer, John Murphy, William Boag, Wei-Hung Weng, Di Jin, Tristan Naumann, and Matthew McDermott. 2019. Publicly available clinical BERT embeddings. In *Proceedings of the 2nd Clinical Natural Language Processing Workshop*, pages 72–78, Minneapolis, Minnesota, USA. Association for Computational Linguistics.

Luis A. Contreras. 2018. An application of natural language processing for triangulation of cognitive load assessments in third level education. Master's thesis, Technological University of Dublin.

Scott H Fraundorf, Zachary A Caddick, Timothy J Nokes-Malach, and Benjamin M Rottman. 2023. Cognitive perspectives on maintaining physicians' medical expertise: IV. best practices and open questions in using testing to enhance learning and retention. *Cogn. Res. Princ. Implic.*, 8(1):53.

Sarah E Giron, Deniz Dishman, Susan P McMullan, Jared Riel, Timothy Newcomer, Dennis Spence, and Shahid A Choudhry. 2021. Longitudinal assessment: A strategy to improve continuing professional certification. *J. Prof. Nurs.*, 37(6):1140–1148.

Le An Ha, Victoria Yaneva, Peter Baldwin, and Janet Mee. 2019. Predicting the difficulty of multiple choice questions in a high-stakes medical exam. In *Proceedings of the Fourteenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 11–20, Florence, Italy. Association for Computational Linguistics.

Marian Hickendorff, Willem J Heiser, Cornelis M van Putten, and Norman D Verhelst. 2009. How to measure and explain achievement change in large-scale assessments: A rejoinder. *Psychometrika*, 74(2):367–374.

Enshuo Hsu and Kirk Roberts. 2025. LLM-IE: a python package for biomedical generative information extraction with large language models. *JAMIA Open*, 8(2):ooaf012.

Sunbeom Kwon and Susu Zhang. 2025. Explaining performance gaps with problem-solving process data via latent class mediation analysis. *Psychometrika*, 90(5):1–29.

Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2020. Biobert: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*, 36(4):1234–1240.

Ming Li, Hong Jiao, Tianyi Zhou, Nan Zhang, Sydney Peters, and Robert W Lissitz. 2025. Item difficulty modeling using fine-tuned small and large language models. *Educ. Psychol. Meas.*, 85(6):1065–1090.

Ying Liu and Steven Culpepper. 2025. Designing learning intervention studies: Identifiability of heterogeneous hidden markov models. *Psychometrika*, 90(4):1258–1283.

Bengt O Muthen and Siek-Toon Khoo. 1998. Longitudinal studies of achievement growth using latent variable modeling. *Learn. Individ. Differ.*, 10(2):73–101.

Laila Rasmy, Yang Xiang, Ziqian Xie, Cui Tao, and Degui Zhi. 2021. Med-BERT: pretrained contextualized embeddings on large-scale structured electronic health records for disease prediction. *NPJ Digit. Med.*, 4(1):86.

Henry L Roediger and Jeffrey D Karpicke. 2006. Test-enhanced learning: taking memory tests improves long-term retention. *Psychol. Sci.*, 17(3):249–255.

Daniel P Schmidt. 2020. Identifying knowledge gaps using a graph-based knowledge representation. Master's thesis, Wright State University.

Matthew Shardlow, Fernando Alva-Manchego, Riza Batista-Navarro, Stefan Bott, Saul Calderon Ramirez, Rémi Cardon, Thomas François, Akio Hayakawa, Andrea Horbach, Anna Hülsing, Yusuke Ide, Joseph Marvin Imperial, Adam Nohejl, Kai North, Laura Occhipinti, Nelson Peréz Rojas, Nishat Raihan, Tharindu Ranasinghe, Martin Solis Salazar, and 3 others. 2024. The BEA 2024 shared task on the multilingual lexical simplification pipeline. In *Proceedings of the 19th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2024)*, pages 571–589, Mexico City, Mexico. Association for Computational Linguistics.

David Spiegelhalter, Christopher Sherlaw-Johnson, Martin Bardsley, Ian Blunt, Christopher Wood, and Olivia Grigg. 2012. Statistical methods for healthcare regulation: Rating, screening and surveillance. *J. R. Stat. Soc. Ser. A Stat. Soc.*, 175(1):1–47.

David J Spiegelhalter. 2005. Funnel plots for comparing institutional performance. *Stat. Med.*, 24(8):1185–1202.

Matthias von Davier, Xueli Xu, and Claus H Carstensen. 2011. Measuring growth in a longitudinal large-scale assessment with a general latent variable model. *Psychometrika*, 76(2):318–336.

Robyn C Ward, Kathy A Baker, Dennis Spence, Colleen Leonard, Alysha Sapp, and Shahid A Choudhry. 2023. Longitudinal assessment to evaluate continued certification and lifelong learning in healthcare professionals: A scoping review. *Eval. Health Prof.*, 46(3):199–212.

Isabella Catharina Wiest, Fabian Wolf, Marie-Elisabeth Leßmann, Marko van Treeck, Dyke Ferber, Jiefu Zhu, Heiko Boehme, Keno K Bressem, Hannes Ulrich, Matthias P Ebert, and Jakob Nikolas Kather. 2025. A software pipeline for medical information extraction with large language models, open source and suitable for oncology. *NPJ Precis. Oncol.*, 9(1):313.

Kazuhiro Yamaguchi, Haruhiko Mitsunaga, Shun Saso, and Yuri Uesaka. 2026. Predicting multilevel growth trajectories using a random-effect diagnostic classification model. *Psychometrika*, pages 1–46.

How Much Training Data Is Enough? Fine-Tuning LLMs for Short-Answer Scoring

Joshua A. McGrane
University of Melbourne
joshua.mcgrane@unimelb.edu.au

David Torres Irribarra
Pontificia Universidad Católica de Chile
davidtorres@uc.cl

Abstract

Fine-tuned models approached state-of-the-art agreement using small labelled sets. Open-weight models met operational criteria on seven of ten items, with a median of 50 responses per score point among passing items. A single marking exercise may supply enough data for automated short-answer scoring.

1 Introduction

Automated short-answer grading (ASAG) has a long history, but instruction-tuned large language models (LLMs) have reset the practical questions that assessment bodies face. Recent work shows that LLMs can approximate human scores on constrained short-answer items, whether prompted with a rubric and a small number of exemplars (Jiang and Bosch, 2024; Lee et al., 2024; Huang and Wilson, 2025) or fine-tuned on human-scored responses (Latif and Zhai, 2024; Ormerod and Kwako, 2024; Ormerod and Kehat, 2025). What it does not establish is whether such systems meet the agreement criteria that high-stakes assessment already requires of automated scoring, or how many human-scored responses are needed to meet those criteria.

Two reporting practices limit the available evidence. The first is that agreement is often reported primarily as quadratic weighted kappa (QWK). The evaluation frameworks developed within testing organisations, notably those of Williamson et al. (2012) and Loukina et al. (2020), require more than chance-corrected agreement with a single rater; machine and human mean scores must be sufficiently close, and machine-human agreement must not fall too far below that of two trained raters on the same item. QWK alone does not establish either requirement. QWK-only reporting can therefore overstate operational readiness, as the results below illustrate. The second is that fine-tuning studies train on whatever labelled data the benchmark supplies, typically a thousand or more responses per item, which tells a practitioner what a large annotation effort achieves but not how agreement grows with the size of the labelled set, nor whether the choice of model matters once fine-tuning begins. Ormerod and Kwako (2024) illustrate both, as they fine-tuned four open-weight models on fixed ASAP-SAS training sets and reported QWK alone.

This paper treats the number of labelled responses as an explicit experimental factor and evaluates fine-tuning against the operational criteria of educational measurement. We fine-tuned GPT-4o-mini, as a closed-source comparator, and gpt-oss-20b and gpt-oss-120b, as small and large open-weight models, on the ten ASAP-SAS items at six training set sizes from 10 to 100 labelled responses per score point.[1] Four questions organise the analysis. First, how many labelled responses per score point are needed before a fine-tuned model satisfies a defensible set of operational criteria? Second, once fine-tuned, do closed-source, small open-weight and large open-weight models still differ? Third, can an item's operational readiness, and which agreement criterion it fails, be anticipated from its rubric and score distribution before collecting additional labelled responses? Fourth, how does agreement reached with labelled sets of this size compare with that of the published systems for the same items?

2 Data

2.1 The ASAP-SAS benchmark

The Automated Student Assessment Prize Short Answer Scoring (ASAP-SAS) dataset (Shermis, 2015) comprises ten short-answer items from standardised assessments in United States schools: three science items (1, 2 and 10), two biology items (5 and 6) and five English Language Arts (ELA)

[1] The companion prompting experiment on the same items and models is reported elsewhere.

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 306–314
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

items (3, 4, 7, 8 and 9). Items 1, 2, 5 and 6 are scored 0 to 3 and the remainder 0 to 2 (Table 1).[2] The two biology items are extremely skewed, with 78% and 85% of responses at the lowest score point and fewer than 3% at the highest; the Science and ELA items are comparatively balanced.

2.2 Evaluation set and training pool

For each item, 600 responses were randomly selected from the public training release, stratified by score to reproduce its score-point distribution, and held out as the evaluation set for every model at every training set size. The evaluation set was drawn from the training release rather than the public test split because the training release is the only part of the dataset in which every response carries two independent human scores. Two of the criteria of Section 3.4 cannot be computed without the second rating: the degradation criterion requires human-human QWK on the same responses as machine-human QWK, and PRMSE requires an estimate of the rater error variance. The remaining responses (658 to 1,193 per item) form the pool from which fine-tuning examples were drawn, less those used as in-context exemplars in a companion experiment. There is no validation set because no hyperparameter was tuned, i.e., the procedure of Section 3.2 was fixed before the first run.

Agreement statistics were computed against the first rater. Loukina et al. (2020) recommend at least 1,000 double-scored responses for PRMSE, but suggest that 500 may suffice when the inter-rater correlation exceeds 0.65; every item meets these latter conditions.

3 Method

3.1 Models

Three models were fine-tuned. GPT-4o-mini (gpt-4o-mini-2024-07-18) is closed-source and was fine-tuned through the OpenAI fine-tuning API.[3] gpt-oss-20b and gpt-oss-120b are open-weight mixture-of-experts models released by OpenAI in 2025, with approximately 21 and 117 billion parameters. They were selected to compare closed-source and open-weight models from one developer, including two open-weight models of different sizes.

3.2 Fine-tuning procedure

Each model was fine-tuned at six training set sizes, $s \in \{10, 20, 30, 50, 70, 100\}$, where s is the target number of labelled training responses per score point on each item. At each s a single adapter (for GPT-4o-mini, a single fine-tuned model) was trained on the pooled training sets of all ten items, each response carrying its own item's system message (Section 3.3), and that model was evaluated on every item. We fixed s per score point rather than in total because it is the sparsely populated upper score points, not the modal one, that limit agreement on skewed items. Training responses were drawn by score-stratified random sampling with a fixed seed, with nested draws so that each training set contained the next smaller set.

The open-weight models were fine-tuned with rank-stabilised LoRA adapters (Hu et al., 2022; Kalajdzievski, 2023) on the attention projections (rank 8, alpha 16), for three epochs at a learning rate of 2×10^{-4} and an effective batch size of 16, on a single 80 GB A100. gpt-oss-20b was trained with its base weights in bfloat16; gpt-oss-120b does not fit at that precision, so its base weights were quantised to 4 bits and frozen in the manner of QLoRA (Dettmers et al., 2023), with the adapter in full precision. Loss was computed on the completion tokens only, matching the default of the OpenAI API, through which GPT-4o-mini was fine-tuned for two epochs with otherwise default hyperparameters. The base model at $s = 0$ was the zero-shot reference.

3.3 Prompt construction and training examples

Each training example, and each evaluation prompt, consisted of a system message and a user message. The system message was written for each item as a translation of the published ASAP-SAS rubric into explicit scoring instructions and was held fixed across models and training set sizes. It comprised a role statement (an automated scorer with 'knowledge' of the subject matter), a task statement giving the item stem, any source materials and the list of correct response elements, with explicit exclusions and a note that students may express an element in different words, and a scoring statement mapping the number of correct elements to score points and requiring a single integer in reply. The stems, ma-

[2] The released rubrics for the items marked 'key elements' enumerate a set of required elements and map the number present to a score point; we retain the dataset's own term. Holistic rubrics call for a single overall judgement of the response.

[3] GPT-4o itself was not fine-tuned because API pricing made a sweep over training set sizes prohibitive.

Item	Domain	Scale	Rubric style	Human-human QWK [95% CI]	Training pool
1	Science	0–3	key elements	0.949 [0.939, 0.957]	1072
2	Science	0–3	key elements	0.915 [0.901, 0.930]	658
3	ELA	0–2	holistic	0.761 [0.735, 0.784]	1193
4	ELA	0–2	holistic	0.749 [0.719, 0.777]	1042
5	Biology	0–3	key elements	0.955 [0.941, 0.965]	1175
6	Biology	0–3	key elements	0.962 [0.944, 0.975]	1177
7	ELA	0–2	key elements	0.971 [0.964, 0.978]	1184
8	ELA	0–2	holistic	0.862 [0.842, 0.879]	1184
9	ELA	0–2	holistic	0.833 [0.815, 0.851]	1183
10	Science	0–2	key elements	0.884 [0.867, 0.900]	1025

Table 1: ASAP-SAS items, rubric style, human-human agreement and training pool size. Human-human QWK is computed from the two independent raters on the full public training release, with 2,000-sample bootstrap intervals. Training pool is the number of responses remaining after the 600 evaluation responses are removed.

terials and rubrics are those released in 2012; the system messages that paraphrase them are ours and available on request. The user message contained the student response. In the fine-tuning files the assistant turn was the human score; at evaluation it was generated at temperature 0.

No scored exemplar appeared in any prompt at any training set size; the labelled responses entered the model only through gradient updates, and the evaluation prompt was identical across all s, including $s = 0$. The zero-shot condition was therefore a rubric-only baseline in the strict sense, and the fine-tuned conditions differed from it only in the trained weights. For the gpt-oss models the 'reasoning' channel was suppressed at training and inference so that the completion was the bare score.

3.4 Agreement criteria

Every (model, item, s) condition was evaluated on the 600 evaluation responses against the criteria of Williamson et al. (2012), which must be met simultaneously: QWK against the first human rater of at least 0.70; degradation of QWK relative to human-human QWK on the same item of no more than 0.10; a standardised mean difference (SMD) between machine and human scores, using the pooled standard deviation, of magnitude no greater than 0.15.

Alongside these we report the proportional reduction in mean squared error (PRMSE) of Loukina et al. (2020). We use 0.70 as a minimum reference level, consistent with McCaffrey et al. (2022) and Casabianca et al. (2023). Higher thresholds may be needed depending on assessment stakes and the contribution of automated scores to reported results; McCaffrey et al. suggest 0.95 when automated scores must be interchangeable with human ratings. Let h_1 and h_2 be the two human scores, m the machine score, $\bar{h} = (h_1 + h_2)/2$ and $n = 600$. Then

$$\begin{aligned} \sigma_e^2 &= \sum (h_1 - h_2)^2 / 2n, \\ \sigma_\tau^2 &= \mathrm{Var}(\bar{h}) - \sigma_e^2/2, \\ \mathrm{MSE} &= \mathrm{mean}\left((m - \bar{h})^2\right) - \sigma_e^2/2, \\ \mathrm{PRMSE} &= 1 - \mathrm{MSE}/\sigma_\tau^2. \end{aligned}$$

PRMSE measures the reduction in prediction error relative to predicting the mean true score, after adjusting for rater error, and is sensitive to differences in the location and scale of machine scores. It is a separate criterion, not a fourth component of the Williamson et al. set.

3.5 Statistical analysis

To model agreement as a function of training set size we fitted a linear mixed model to the Fisher-z transform of QWK in each condition, $z = \mathrm{atanh}(\mathrm{QWK})$, with each condition weighted by the inverse of a bootstrap variance estimate on the z scale from 1,000 resamples of the evaluation responses. Because a single log-linear curve through the zero-shot and fine-tuned conditions would misspecify the zero-shot point as the limit of a curve whose slope is identified only on the fine-tuned side, we fitted a piecewise specification with a separate term for the zero-shot condition,

$$\begin{aligned} z \sim\ & \mathrm{zs} + \mathrm{model} + \mathrm{zs{:}model} \\ & + \mathrm{model{:}}\log(1 + s_{\mathrm{ft}}) \\ & + \left(1 + \log(1 + s_{\mathrm{ft}}) \mid \mathrm{item}\right), \end{aligned}$$

where zs indicates the zero-shot condition and s_{ft} equals s for fine-tuned conditions and zero otherwise. The piecewise specification was preferred to a single log-linear curve by a wide margin ($\Delta\mathrm{AIC} = -40.7$, $\Delta\mathrm{BIC} = -24.0$). The

model was fitted by REML in lme4 (Bates et al., 2015), with Satterthwaite degrees of freedom from lmerTest (Kuznetsova et al., 2017) and Tukey-adjusted pairwise contrasts on the fine-tuning slopes. The same fixed and random effects were refitted to PRMSE, degradation and |SMD| on their original scales without weights. The QWK and PRMSE slope contrasts and the QWK zero-shot model contrasts were also evaluated by an item-level cluster bootstrap ($B = 500$). With only ten items, we deferred to the bootstrap when it disagreed with model-based inference. The degradation and |SMD| fits describe population curves, with model-based confidence intervals. Item-level results are reported as the smallest observed s at which a (model, item) condition satisfies all three Williamson et al. criteria, and as a model-based estimate from a log-linear fit to each (model, item) series.

4 Results

4.1 Agreement as a function of training set size

Figure 1 plots the marginal curves for the four criteria against s for the three models, as fixed-effect estimates from the piecewise model; Table 2 gives fitted QWK across training set sizes and all four outcomes at $s = 100$. Four features warrant comment. First, agreement rose steeply with few labelled responses. From a zero-shot baseline near 0.60, ten responses per score point raised marginal QWK to approximately 0.68, twenty carried it above the 0.70 floor, and one hundred brought it to approximately 0.79, against a mean human-human QWK of 0.888. The gain was steepest over the first ten to fifty responses.

Second, the zero-shot condition was not the origin of the fine-tuning curve. The zero-shot main effect was significant ($F(1, 184.5) = 16.9$, $p < 0.001$), and each pretrained baseline sat above the value that back-extrapolating the fine-tuning curve to $s = 0$ would predict, by 0.18, 0.13 and 0.07 on the z scale for GPT-4o-mini, gpt-oss-20b and gpt-oss-120b ($p < 0.001$, 0.016 and 0.18). By $s = 10$ every model already exceeded its zero-shot value, so whether agreement dips between $s = 0$ and $s = 10$ or the curve simply flattens below the smallest training set size cannot be told from these data. The displacement did not differ by model ($p = 0.35$ for the zero-shot by model interaction).

Third, marginal PRMSE tracked QWK but sat lower: at $s = 100$ the refit gave 0.635, 0.641 and 0.655, with upper bootstrap limits of 0.72 to 0.74. The average item did not reach the 0.70 threshold for any training set size (Section 5.3).

s	GPT-4o-mini	gpt-oss-20b	gpt-oss-120b
QWK by training set size			
0 (zero-shot)	0.620	0.599	0.567
10	0.676	0.681	0.687
20	0.714	0.717	0.724
30	0.735	0.738	0.744
50	0.760	0.762	0.768
70	0.776	0.777	0.783
100	0.791	0.792	0.798
Other outcomes at $s = 100$			
PRMSE	0.635	0.641	0.655
Degradation	−0.092	−0.092	−0.087
|SMD|	0.096	0.079	0.077

Table 2: Marginal agreement from the fixed effects of the four piecewise models. QWK is back-transformed from the Fisher-z scale; the other outcomes are fitted without weights on their original scales. The zero-shot estimate is separate from the fine-tuning curve. Figure 1 shows confidence intervals.

Fourth, fitted degradation first met the -0.10 floor at $s = 100$ among the observed training sizes for all three models, making it the binding population criterion. For |SMD|, GPT-4o-mini first met the 0.15 ceiling at $s = 50$, whereas both open-weight models met it at every observed fine-tuned size. These comparisons concern fitted means; the confidence intervals in Figure 1 show the uncertainty around threshold crossings.

4.2 Comparison of fine-tuning learning rates

The estimated learning rates were similar across models, with no statistically detectable pairwise differences. The fine-tuning log-slopes on the z scale were 0.114 (SE 0.016) for GPT-4o-mini, 0.110 (0.016) for gpt-oss-20b and 0.112 (0.016) for gpt-oss-120b. No pairwise difference exceeded 0.005, all three Tukey-adjusted contrasts returned $p > 0.97$, and the cluster bootstrap was consistent (all pairwise $p_{\text{boot}} > 0.81$). The PRMSE refit showed the same pattern, with slopes of 0.162, 0.152 and 0.151 and all pairwise $p_{\text{boot}} > 0.67$. The zero-shot baselines differed in point estimate, GPT-4o-mini sitting 0.08 above gpt-oss-120b on the z scale (Tukey $p = 0.015$), but the contrast was not significant using cluster-robust inference ($p_{\text{boot}} = 0.29$).

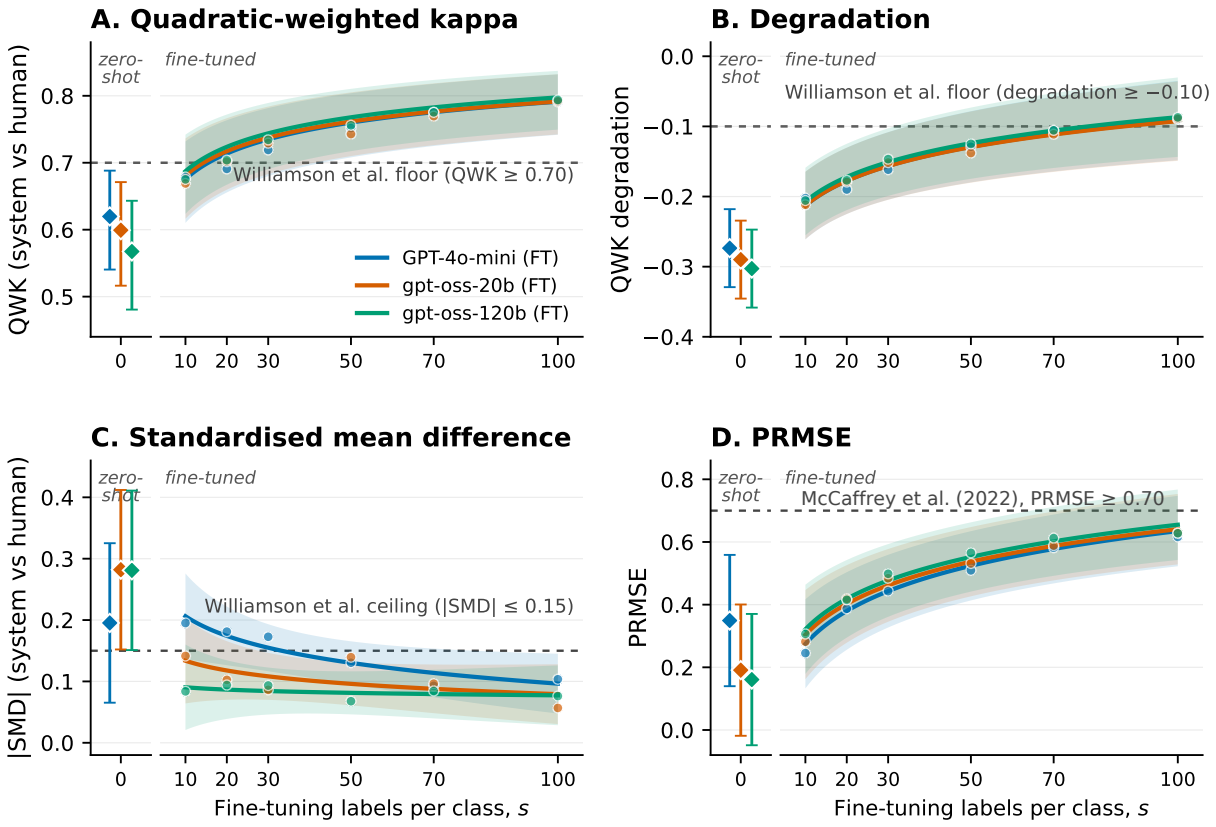


Figure 1: Marginal curves for the four criteria against labelled responses per score point: QWK (A), degradation of QWK below the human-human QWK (B), absolute standardised mean difference (C) and PRMSE (D), from the piecewise mixed model fitted to each outcome (QWK on the Fisher-z scale and back-transformed), with 95% confidence ribbons from the fixed effects. Zero-shot estimates are plotted on a separate axis at $s = 0$, disconnected from the fine-tuning curves, which are drawn over the observed training set sizes only. Points are the means of the ten observed item conditions at each training set size for each model. Dashed rules mark the Williamson et al. thresholds (QWK ≥ 0.70, degradation ≥ -0.10, $|\text{SMD}| \leq 0.15$) and the minimum PRMSE reference level of 0.70 (McCaffrey et al., 2022; Casabianca et al., 2023); its suitability depends on the intended use (Section 3.4).

4.3 Item-level results against the operational criteria

Figure 2 shows, for each of the 30 (model, item) conditions, the smallest s at which PRMSE reaches 0.70, with a marker where the three Williamson et al. criteria are satisfied; Table 3 gives the smallest observed s at which those three criteria are met. Within $s \leq 100$, gpt-oss-20b and gpt-oss-120b each satisfied all three on seven of ten items (1, 2, 3, 4, 6, 9 and 10), and GPT-4o-mini on five (1, 2, 6, 9 and 10). Across the passing items the median observed s was 50 for all three models; the model-based estimate for gpt-oss-20b was 42.

GPT-4o-mini's two additional failures, on items 3 and 4, were failures of standardised mean difference: its QWK crossed the 0.70 floor, but the fine-tuned model shifted the score distribution relative to the human raters in a way that the open-weight models, at comparable QWK, did not.

4.4 Items that fail the criteria

Three features help explain whether items met the criteria: agreement generally rose with the number of labelled responses; high human-human agreement imposed a demanding degradation threshold; and extreme score skew limited the available examples at upper score points. These constraints can overlap, as item 6 illustrates. Items 3, 4, 7 and 8 generally had lower QWK across the published systems in Table 4, suggesting recurring item-level difficulties across scorers.

Items 5, 7 and 8 failed for every model at every training set size, for different reasons. Item 5 has an enumerated rubric and an evaluation-set human-human QWK of 0.951, so the degradation criterion requires machine-human QWK of at least 0.851. Fine-tuned agreement exceeded the absolute 0.70 floor but failed this stricter threshold. Its two upper score points also have the fewest available

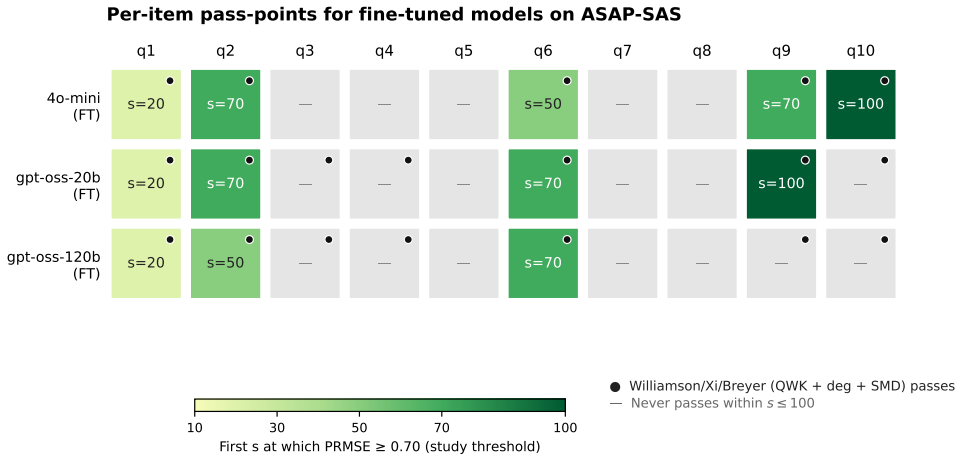


Figure 2: Smallest training set size at which PRMSE reaches 0.70 for each model and item (cell shading), with a marker where the three Williamson, Xi and Breyer criteria are satisfied. Grey cells do not reach the criterion within $s \leq 100$.

Item	4o-mini	oss-20b	oss-120b	Binding criterion
1	70	50	30	degradation
2	50	50	50	degradation
3	–	70	50	absolute 0.70
4	–	50	70	absolute 0.70
5	–	–	–	degradation
6	50	100	100	degradation
7	–	–	–	degradation
8	–	–	–	degradation
9	20	10	30	degradation
10	30	70	100	degradation

Table 3: Smallest observed training set size (responses per score point) at which all three Williamson et al. criteria are satisfied, by item and model, and the criterion last satisfied (binding) on each item for the models that pass; for GPT-4o-mini on items 3 and 4 the binding criterion is the SMD. A dash indicates that the criteria are not met within $s \leq 100$.

responses (25 and 19), and the confusion matrices showed almost exclusively under-prediction of these rare high scores. Item 6 shares the skew, has more upper-score responses (44 and 30), and passed at $s = 50$ for GPT-4o-mini and $s = 100$ for the open-weight models. Item 7 also has high evaluation-set human-human QWK (0.964), and Jiang and Bosch (2024) report over-scoring under GPT-4 prompting; the criterion is demanding and the item appears sensitive to how the rubric is specified. Item 8 has a lower machine-human QWK target (0.771) but a holistic rubric on which Jiang and Bosch document machine scores 'distributed across the entire range'. The fine-tuned models showed similar but less pronounced score dispersion. These results warrant examining the rubric and retaining human scoring while agreement remains inadequate; they do not establish that further labelling would be ineffective.

4.5 Rubric structure as an item-level moderator

Rubric style accounted for 47% of the estimated item random-intercept variance. Adding the binary distinction in Table 1 (key elements versus holistic) as a main effect produced this reduction (likelihood-ratio $\chi^2(1) = 7.25$, $p = 0.007$). Allowing rubric style to moderate the slope did not improve fit ($\chi^2(1) = 0.28$, $p = 0.60$). Items with enumerated rubrics had higher estimated QWK throughout, with similar estimated gains from additional labelled responses (slopes of 0.106 to 0.111 against 0.116 to 0.120 for holistic items). Rubric style was associated with agreement levels, but did not by itself determine whether items met all three criteria.

4.6 Comparison with published systems

Fine-tuning reached near state-of-the-art agreement with substantially fewer labelled responses per item (Table 4). At $s = 100$, mean observed QWK was 0.79, within 0.01 of the 0.796 ensemble benchmark identified as state of the art by Ormerod and Kwako (2024). That benchmark comes from the pretrained-model ensemble of Ormerod (2022). Our models used at most 300 or 400 labelled responses per item, depending on the score scale, compared with 1,022 to 1,446 training responses plus development data in the ensemble study. At $s = 50$, with at most 150 or 200 responses per item, mean QWK

was 0.75, matching BERT-base and the average of the four fine-tuned open-weight LLMs reported by Ormerod and Kwako (2024). Their generative models used 984 to 1,278 training responses per item. The $s = 100$ mean also exceeded the competition average of 0.73 calculated from Shermis (2015).

A published small-sample comparison used an n-gram support vector machine trained on 270 responses per item, which reached 0.64, 0.56 and 0.64 on items 1, 2 and 10 (Horbach et al., 2018); at $s = 20$, with at most 80 labelled responses per item, our fine-tuned models reached 0.84, 0.76 and 0.75 on the same items.

These counts are per-item labelling requirements for our model trained jointly on all ten items. Different evaluation samples and reference scores make the comparisons indicative, rather than evidence of superiority or equivalence under a common test. Our results are against the first human rater; the public test set used by BERT-base and the ensemble supplies one reference score.

5 Discussion

5.1 The size of the labelled set

On this benchmark, passing items required a median of fifty labelled responses per score point to satisfy all three Williamson et al. criteria, and none required more than one hundred. The comparison studies used roughly 1,000 to 1,400 training responses per item, and a reader of that literature could reasonably conclude that adequate agreement requires annotation on that scale. The present results show that smaller per-item labelled sets can suffice. For an item scored 0 to 3, fifty responses per score point corresponds to 200 labelled responses. A single marking exercise may supply this quantity, bringing fine-tuned scoring within reach of assessment units and research groups with limited annotation resources, subject to item-level evaluation.

5.2 Model selection

The models had similar estimated QWK and PRMSE learning rates; deployment requirements and item-level criterion failures therefore inform model choice. gpt-oss-20b is a practical default: it met the Williamson et al. criteria on the same seven items as gpt-oss-120b, avoided the mean-score bias that caused GPT-4o-mini to fail two ELA items, and can be hosted locally so that student responses remain within the institution. It is the smaller open-weight model and requires fewer local resources than gpt-oss-120b. The similar learning rates are consistent with Ormerod and Kwako (2024), whose four fine-tuned models had mean QWK values spanning 0.749 to 0.759, and with Ormerod and Kehat (2025), whose 130M-parameter Mamba and 8B-parameter Llama models reached QWK of 0.797 and 0.792 on ASAP 2.0 essay scoring. These comparisons support choosing a smaller model when it meets the required item-level criteria; they do not establish equivalence across all scoring outcomes.

5.3 The limitations of QWK-only reporting

Two results illustrate the limits of QWK-only reporting. First, acceptable QWK can coexist with excessive mean-score bias. On items 3 and 4, fine-tuned GPT-4o-mini cleared the 0.70 floor and would meet that QWK-only threshold, but its scores were shifted relative to the human raters by more than the SMD criterion allows. The second is item-level heterogeneity, which a single benchmark mean conceals. At $s = 100$, the per-item QWK in Table 4 ranged from 0.71 to 0.90 around a mean of 0.79, and three items failed the operational criteria at every training set size; PRMSE showed the converse, with the curve for the average item never reaching 0.70 while seven of the ten items reached it individually. A benchmark mean therefore obscures variation in whether individual items meet the operational criteria.

5.4 Benchmark contamination

ASAP-SAS has been public since 2012, so its items, rubrics and released responses may have been encountered during pretraining. Three observations qualify the interpretation of the results. First, the zero-shot baselines were poor, failing both the Williamson et al. criteria and the PRMSE threshold on every item. Second, our learning-curve slopes were estimated from fine-tuned conditions that shared the same pretrained model within each model family. This holds pretraining exposure constant within each curve, but does not establish that contamination leaves slopes or estimated training requirements unaffected. Third, the gpt-oss models had the weakest zero-shot performance despite their later release. This provides no clear evidence of effective recall of the released scores, but cannot rule out contamination. Replication on an unreleased item bank is therefore needed to

System	1	2	3	4	5	6	7	8	9	10	Mean
Human raters, this study	0.95	0.91	0.79	0.72	0.95	0.97	0.96	0.87	0.82	0.87	0.88
2012 competition, competitor mean	0.82	0.76	0.71	0.64	0.78	0.81	0.65	0.60	0.78	0.73	0.73
BERT-base, fine-tuned	0.85	0.77	0.69	0.72	0.85	0.84	0.68	0.60	0.83	0.72	0.75
Ensemble, pretrained language models	0.88	0.89	0.72	0.75	0.81	0.82	0.73	0.70	0.87	0.78	0.80
GPT-4, one exemplar per score point	0.72	0.72	0.63	0.52	0.77	0.80	0.50	0.55	0.70	0.87	0.68
Open-weight LLMs, fine-tuned, average of four models	0.87	0.81	0.68	0.64	0.80	0.83	0.73	0.69	0.77	0.71	0.75
This study, zero-shot, average of three models	0.73	0.66	0.54	0.29	0.62	0.77	0.53	0.57	0.51	0.69	0.59
This study, $s = 50$, average of three models	0.85	0.83	0.69	0.70	0.75	0.81	0.68	0.66	0.76	0.76	0.75
This study, $s = 100$, average of three models	0.90	0.85	0.72	0.73	0.81	0.88	0.74	0.71	0.77	0.81	0.79

Table 4: Per-item QWK on the ten ASAP-SAS items. Sources: competition, Shermis (2015); BERT-base and ensemble, Ormerod (2022); open-weight LLMs, Ormerod and Kwako (2024); GPT-4, Jiang and Bosch (2024). Means are arithmetic averages across items; values are rounded to two decimals. Evaluation sets differ: competition results use the private leaderboard; BERT-base and the ensemble use the public test set; the four generative LLMs use the study's train/development/test partition; GPT-4 uses samples of 100 to 356 responses per item; this study uses 600 responses held out from the training release (Section 2.2).

assess whether the reported training requirements generalise.

6 Limitations

Two limitations bear on the item-level estimates. First, the largest training condition targeted 100 labelled responses per score point, but four items had insufficient examples: mildly on items 2 and 4, where one score point is capped at 85 and 77 responses, and severely on items 5 and 6, where increases in training set size mainly added responses at the modal score point, leaving the rare upper score points underrepresented. Second, the public release contains no respondent background information, so the SMD criterion cannot be supplemented by subgroup analyses of the kind Huang and Wilson (2025) report for English language learners, and no fairness claim across such groups is made. Results are from a single English-language benchmark of ten items, and the training set sizes should not be generalised to item banks with different score distributions or rubric styles.

7 Conclusion

The four questions posed in the introduction can be addressed on the ASAP-SAS items. First, each open-weight model satisfied all three Williamson et al. criteria on seven of ten items, and GPT-4o-mini on five. Among passing items, the median was fifty labelled responses per score point, and none required more than one hundred. Second, the models showed similar gains from additional labelled responses, with no detectable differences in learning-curve slopes. Deployment requirements and item-level performance favour gpt-oss-20b among the models examined: it can run locally and met the criteria on the same items as the larger open-weight model. Third, rubric features, score distributions and human-human agreement help explain item-level difficulties, but do not provide a sufficient rule for predicting which items will pass. Fourth, joint training on small labelled sets reached near state-of-the-art agreement: mean QWK of 0.79, within 0.01 of the ensemble benchmark highlighted by Ormerod and Kwako (2024), with fewer labelled responses per item. A set that a single marking exercise may supply can therefore be sufficient for items that meet the operational criteria. Evaluation against the full set of criteria is needed to identify those items and determine where human scoring remains necessary.

References

Douglas Bates, Martin Mächler, Ben Bolker, and Steve Walker. 2015. Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1):1–48.

Jodi M. Casabianca, Daniel F. McCaffrey, Matthew S. Johnson, Kathryn L. Ricker-Pedley, Ourania Rotou, and Joseph Martineau. 2023. Exploration of the proportional reduction in mean-squared error for evaluating automated scores. Research Memorandum RM-23-01, Educational Testing Service.

Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2023. QLoRA: Efficient finetuning of quantized LLMs. In *Advances in Neural Information Processing Systems*, volume 36.

Andrea Horbach, Sebastian Stennmanns, and Torsten Zesch. 2018. Cross-lingual content scoring. In *Proceedings of the Thirteenth Workshop on Innovative Use of NLP for Building Educational Applications*,

pages 410–419. Association for Computational Linguistics.

Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*.

Yue Huang and Joshua Wilson. 2025. Evaluating LLM-based automated essay scoring: Accuracy, fairness, and validity. In *Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con): Works in Progress*, pages 71–83. National Council on Measurement in Education.

Lan Jiang and Nigel Bosch. 2024. Short answer scoring with GPT-4. In *Proceedings of the Eleventh ACM Conference on Learning @ Scale (L@S '24)*. ACM.

Damjan Kalajdzievski. 2023. A rank stabilization scaling factor for fine-tuning with LoRA. *Preprint*, arXiv:2312.03732.

Alexandra Kuznetsova, Per B. Brockhoff, and Rune H. B. Christensen. 2017. lmerTest package: Tests in linear mixed effects models. *Journal of Statistical Software*, 82(13):1–26.

Ehsan Latif and Xiaoming Zhai. 2024. Fine-tuning ChatGPT for automatic scoring. *Computers and Education: Artificial Intelligence*, 6:100210.

Gyeong-Geon Lee, Ehsan Latif, Xuansheng Wu, Ninghao Liu, and Xiaoming Zhai. 2024. Applying large language models and chain-of-thought for automatic scoring. *Computers and Education: Artificial Intelligence*, 6:100213.

Anastassia Loukina, Nitin Madnani, Aoife Cahill, Lili Yao, Matthew S. Johnson, Brian Riordan, and Daniel F. McCaffrey. 2020. Using PRMSE to evaluate automated scoring systems in the presence of label noise. In *Proceedings of the Fifteenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 18–29. Association for Computational Linguistics.

Daniel F. McCaffrey, Jodi M. Casabianca, Kathryn L. Ricker-Pedley, René R. Lawless, and Cathy Wendler. 2022. Best practices for constructed-response scoring. *ETS Research Report Series*, 2022(1):1–58.

Christopher Ormerod. 2022. Short-answer scoring with ensembles of pretrained language models. *Preprint*, arXiv:2202.11558.

Christopher Ormerod and Gitit Kehat. 2025. Long context automated essay scoring with language models. In *Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con): Full Papers*, pages 35–42. National Council on Measurement in Education.

Christopher Ormerod and Alexander Kwako. 2024. Automated text scoring in the age of generative AI for the GPU-poor. *Chinese/English Journal of Educational Measurement and Evaluation*, 5(3):Article 5.

Mark D. Shermis. 2015. Contrasting state-of-the-art in the machine scoring of short-form constructed responses. *Educational Assessment*, 20(1):46–65.

David M. Williamson, Xiaoming Xi, and F. Jay Breyer. 2012. A framework for evaluation and use of automated scoring. *Educational Measurement: Issues and Practice*, 31(1):2–13.

Evaluating Machine Learning, Transformer, and LLM Approaches for Large-Scale Automated Essay Scoring

Cyrus Sai-Cheong Chan[1], Charles Anifowose[1], Su Zhang[2], Winnie Wing-Yee Tse[2], Nizam Radwan[2], Hyunah Kim[2]

[1] Vretta
[2] Education Quality and Accountability Office

Abstract

This study compared machine learning, transformer fine-tuning, prompt-based LLM, and novel ensemble approaches for automated essay scoring in a Canadian large-scale provincial assessment. Fine-tuned transformers achieved the highest reliability, followed by machine learning models. Meanwhile, prompt-based LLMs provided greater explainability, and ensemble architectures highlighted opportunities for balancing reliability and explainability.

1 Introduction

1.1 Automated Essay Scoring

Automated essay scoring (AES) has evolved across multiple technical generations, encompassing feature-engineered machine learning (ML) models, fine-tuned transformer architectures, and prompt-based large language models (LLMs). Each approach presents distinct trade-offs regarding scoring reliability, system transparency, and operational feasibility (Atkinson & Palma, 2025; Huang & Wilson, 2025; Li & Ng, 2024; Ramesh & Sanampudi, 2022). Although prompt-based LLM scoring has attracted growing research attention (Chu et al., 2025; Hou et al., 2025; Lee et al, 2024), operational implementations in large-scale K-12 assessments have mainly relied on traditional ML-based engines (Zhang et al., 2026).

1.2 Aim

This study evaluates the operational feasibility of diverse AES approaches within a high-stakes, large-scale Canadian provincial assessment. Specifically, it compares traditional ML, transformer fine-tuning, prompt-based LLMs, and novel ensemble architectures across four key dimensions: scoring reliability, subgroup consistency, explainability, and operational feasibility.

2 Sample

This study utilized student responses and scoring data across three long-writing items from two assessment administrations in 2023-2024. Student responses for long-writing items were evaluated on two scoring scales: Topic Development (TD; 6-point scale) and Conventions (CV; 4-point scale). Generic, scale-specific rubrics were established and applied across writing prompts.

A stratified sampling strategy based on administration window, gender, and English Language Learner (ELL) status ensured the representativeness of the sample. Two operational datasets and an expert-audit dataset were constructed to develop and evaluate the AES approaches.

2.1 Operational Dataset A

Operational Dataset A (Operational A) was created for ML model training, transformer fine-tuning, and ensemble architecture evaluation. Comprising 15,000 student responses, it was split into training, validation, and testing sets using a 70/15/15 ratio. The training set was used to fit the ML models and fine-tune the transformers, but was not required for the ensemble approach as it involved no dedicated training phase. The validation set served to compare candidate models within each approach and select the best-performing configurations. The testing set was reserved for final independent model evaluations.

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 315–321
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

To ensure direct comparability across the four AES approaches on a unified benchmark, all best-performing models, including the prompt-based LLM, were evaluated on the held-out Operational A testing set.

2.2 Operational Dataset B

Operational Dataset B (Operational B) was dedicated exclusively to LLM prompt optimization. Comprising 12,000 student responses, it was independently sampled from the same overarching pool of operational responses as Operational A. Consequently, the two datasets shared a partial overlap of 3,191 samples. Operational B was partitioned into training, validation, and testing sets using a 20/40/40 split. The training set facilitated subject matter expert (SME) feedback collection and iterative prompt optimization, the validation set was used to select the best performing prompt templates, and the testing set provided an independent evaluation of the selected prompt template.

2.3 Expert-Audit Dataset

An expert-audit dataset was established as an authoritative benchmark to ensure model alignment with expert scoring standards. The dataset comprised 404 responses for TD and 387 responses for CV across the three long-writing items, covering all responses used in operational scorer qualifying tests and validity checks. Scores were assigned and validated by expert scorers, reflecting rigorous operational standards. Within this dataset, 21 responses overlapped across the TD and CV scales. Relative to the broader operational datasets, this audit set shared 29 overlapping responses with Operational A and 28 with Operational B.

3 Methodology

Four distinct AES approaches were implemented and evaluated. Within each approach, various model configurations were developed and benchmarked to select optimal candidates.

3.1 Machine Learning Approach

Score prediction was framed as a regression task, with clipping and rounding applied during post-processing to align continuous predictions with valid score points. Feature engineering generated 22 linguistic features for TD and 24 for CV. Three scale-specific ML models were trained on identical feature sets: Ridge Regression (RR), Support Vector Regression (SVR), and XGBoost Regression (XGB).

3.2 Transformer Fine-Tuning Approach

The transformer fine-tuning approach represented a middle ground between feature-engineered ML engines and prompt-driven LLMs. Three pretrained transformers, BERT, DistilBERT, and RoBERTa, were fine-tuned for each scale using uniform hyperparameter configurations: 4 epochs, a learning rate of 1e-5, weight decay of 0.01, and FP16 mixed precision.

3.3 Prompt-Based LLM Approach

The prompt-based approach utilized Llama 3 70B, an open-weight foundation model, configured with zero-shot chain-of-thought prompting. Prompt context included scoring rubrics, scorer training guides, additional scoring instructions, item prompts, and student responses. The model generated both numerical scores and qualitative scoring rationales.

Prompt optimization for the TD scale involved four iterative rounds on non-overlapping subsets of the Operational B training set. SME feedback on discrepant cases was systematically incorporated to refine the additional scoring instructions, yielding four distinct prompt templates.

3.4 Ensemble Approach

The ensemble approach integrated traditional ML with prompt-based LLMs through three mechanisms: linguistic feature injection, reference score injection, and combined injection.

For linguistic feature injection, rubric-relevant linguistic features were aggregated into qualitative performance indicators and injected into the top-performing LLM prompt template as objective evidence to assist model reasoning.

For reference score injection, predictions from top-performing ML models were provided within the prompt to inform, but not dictate, LLM evaluations. The LLM retained autonomy to override reference scores if its internal reasoning diverged from ML predictions.

For combined injection, both qualitative performance indicators and ML-predicted reference scores were integrated into the augmented LLM prompt template.

Dataset	Item	Metric	Human	ML	Transformer	LLM	Ensemble
Operational A	Item 1	QWK	0.76	0.83	0.84	0.67	0.76
		SMD	-0.13	-0.08	-0.03	-0.33	-0.14
		EA	57.82	64.67	65.87	44.27	52.40
	Item 2	QWK	0.74	0.78	0.80	0.64	0.71
		SMD	0.02	-0.03	-0.01	0.15	0.16
		EA	59.74	66.40	66.40	49.33	53.87
	Item 3	QWK	0.75	0.85	0.87	0.69	0.76
		SMD	-0.12	0.01	0.04	-0.14	-0.04
		EA	55.04	66.53	67.47	46.27	50.80
Operational B	Item 1	QWK	0.76	-	-	0.70	-
		SMD	-0.13	-	-	-0.25	-
		EA	57.82	-	-	49.00	-
	Item 2	QWK	0.74	-	-	0.67	-
		SMD	0.02	-	-	0.12	-
		EA	59.74	-	-	47.69	-
	Item 3	QWK	0.75	-	-	0.75	-
		SMD	-0.12	-	-	-0.10	-
		EA	55.04	-	-	47.94	-
Expert-audit	Item 1	QWK	0.97	0.97	0.98	0.88	0.92
		SMD	-0.02	-0.06	-0.02	-0.04	-0.13
		EA	82.61	83.96	85.85	50.94	59.43
	Item 2	QWK	0.95	0.94	0.97	0.83	0.90
		SMD	-0.03	-0.13	-0.06	0.14	-0.16
		EA	81.03	76.19	86.39	44.90	55.78
	Item 3	QWK	0.96	0.96	0.98	0.89	0.92
		SMD	-0.05	-0.11	-0.03	0.05	-0.14
		EA	84.23	78.15	86.09	51.66	52.98

Table 1: Model testing results with human-human agreement baseline on Topic Development.

3.5 Evaluation Framework

Human-human agreement baselines were established to benchmark automated model performance. The operational baseline was calculated using all available double-scored operational responses rather than dataset-specific samples, thereby reflecting overall human reliability and providing a common reference benchmark for both Operational Datasets A and B. For the expert-audit dataset, the baseline was calculated by comparing human scores against expert scores, using bootstrapping with 10,000 iterations to derive stable mean agreement statistics across multi-scored responses.

Automated model performance was evaluated relative to the human-human agreement baselines using industry-standard operational metrics with acceptance thresholds based on cross-jurisdictional practices (Zhang et al., 2026):

- Quadratic Weighted Kappa (QWK) ≥ 0.70
- Absolute value of Standardized Mean Difference (SMD) ≤ 0.15
- Exact Agreement (EA) rate ≥ 65%
- QWK degradation from the human-human baseline ≤ 0.10
- EA degradation from the human-human baseline ≤ 5.25 percentage points (pp)

Subgroup performance was monitored across gender and ELL status to evaluate scoring equity. Subgroup analyses were limited to Operational Datasets A and B due to sample size constraints in the expert-audit dataset.

4 Results

Among the trained ML models, XGB achieved the highest scoring performance on TD, while SVR performed best on CV. DistilBERT emerged as the top-performing fine-tuned transformer model across both scales. For prompt-based LLMs, four rounds of prompt optimization were completed on

Dataset	Item	Metric	Human	ML	Transformer	LLM	Ensemble
Operational A	Item 1	QWK	0.61	0.70	0.75	-	-
		SMD	-0.01	0.05	0.14	-	-
		EA	54.38	55.33	61.07	-	-
	Item 2	QWK	0.62	0.63	0.67	-	-
		SMD	0.03	0.13	0.18	-	-
		EA	56.21	54.00	55.60	-	-
	Item 3	QWK	0.60	0.67	0.70	-	-
		SMD	-0.15	0.10	0.16	-	-
		EA	52.05	55.07	56.67	-	-
Operational B	Item 1	QWK	0.61	-	-	0.54	-
		SMD	-0.01	-	-	0.24	-
		EA	54.38	-	-	51.50	-
	Item 2	QWK	0.62	-	-	0.40	-
		SMD	0.03	-	-	0.23	-
		EA	56.21	-	-	46.00	-
	Item 3	QWK	0.60	-	-	0.48	-
		SMD	-0.15	-	-	0.22	-
		EA	52.05	-	-	49.50	-
Expert-audit	Item 1	QWK	0.90	0.88	0.92	0.65	0.72
		SMD	-0.05	0.03	0.03	-0.03	0.02
		EA	77.79	72.26	79.56	45.26	48.91
	Item 2	QWK	0.89	0.85	0.87	0.70	0.79
		SMD	-0.05	0.05	0.12	0.29	0.19
		EA	77.86	70.77	71.54	54.62	58.46
	Item 3	QWK	0.94	0.89	0.89	0.78	0.85
		SMD	0.00	0.13	0.18	0.11	0.17
		EA	85.41	74.17	76.67	61.67	71.67

Table 2: Model testing results with human-human agreement baseline on Conventions.

TD with the final prompt template version producing the strongest performance, whereas the CV prompt remained unoptimized. Among ensemble strategies, reference score injection achieved the highest scoring agreement on TD, while combined injection performed best on CV.

Table 1 and Table 2 present testing results for the top-performing models across all AES approaches on the TD and CV scales, respectively, alongside the human-human baseline. The top-performing ML models demonstrated strong, stable predictive performance. Fine-tuned transformers achieved even slightly higher reliability while eliminating manual feature engineering. Prompt-based LLMs generated human-readable scoring rationales, provided an auditable prompt optimization process, and incorporated complex scoring standards via natural-language instructions. The ensemble approach outperformed standalone prompt-based LLMs while achieving slightly lower reliability than pure ML models, successfully preserving rationale generation capabilities.

Table 3 and Table 4 summarize subgroup performance on the testing sets. ML and transformer models exhibited consistent performance across Female, Male, and ELL subgroups, though performance dropped slightly on the non-ELL subgroup. Prompt-based LLM and ensemble approaches showed lower consistency across Female and Non-ELL subgroups, exhibiting a pronounced drop on the latter.

Prompt-based LLM and ensemble results on CV are omitted from Operational A in Table 2 and Table 4 because no prompt optimization was conducted on CV, leaving no optimized template to evaluate or integrate into the ensemble architecture for that dataset. The LLM results reported for CV on Operational B represent the performance of the unoptimized initial prompt template. Ensemble injection methods for CV were evaluated and compared exclusively on the expert-audit dataset.

Dataset	Subgroup	Metric	ML	Transformer	LLM	Ensemble
Operational A	Female	QWK	0.81	0.82	0.61	0.70
		SMD	-0.07	-0.01	-0.11	-0.01
		EA	67.50	66.96	47.87	52.75
	Male	QWK	0.83	0.85	0.71	0.77
		SMD	0.00	0.01	-0.12	-0.01
		EA	64.23	66.19	45.37	51.96
	ELL	QWK	0.85	0.87	0.72	0.79
		SMD	-0.01	-0.01	-0.13	-0.04
		EA	65.10	66.34	44.58	52.04
	Non-ELL	QWK	0.76	0.78	0.54	0.63
		SMD	-0.07	0.02	-0.09	0.03
		EA	66.64	66.81	48.67	52.67
Operational B	Female	QWK	-	-	0.68	-
		SMD	-	-	-0.09	-
		EA	-	-	49.24	-
	Male	QWK	-	-	0.72	-
		SMD	-	-	-0.08	-
		EA	-	-	48.16	-
	ELL	QWK	-	-	0.73	-
		SMD	-	-	-0.11	-
		EA	-	-	47.43	-
	Non-ELL	QWK	-	-	0.65	-
		SMD	-	-	-0.05	-
		EA	-	-	49.92	-

Table 3: Model subgroup performance on Topic Development.

5 Discussion

These findings illustrate distinct trade-offs among AES approaches in K-12 large-scale assessment. Fine-tuned transformers achieved the highest overall reliability and subgroup consistency, followed by ML models; both met most operational acceptance thresholds despite offering limited transparency into their scoring decisions. Conversely, the open-weight prompt-based LLM yielded lower scoring agreement and met fewer quantitative thresholds, but generated human-readable rationales that provided useful supporting information for SME review and scoring decisions. These findings align with empirical studies showing that fine-tuned transformers outperform traditional regression-based methods in scoring accuracy (Wang et al., 2022; Yang et al., 2020), whereas prompt-engineered LLMs consistently underperform relative to ML and transformer baselines (Stahl et al., 2024; Yancey et al., 2023).

Despite strong natural language understanding capabilities, prompt engineering alone proved insufficient to match the predictive performance of traditional ML and fine-tuned transformers, which is consistent with broader literature (Atkinson & Palma, 2025; Li & Ng, 2024; Ramesh & Sanampudi, 2022; Xiao et al., 2024).

However, the ensemble findings demonstrate that these AES approaches can be mutually reinforcing. Integrating ML predictions into prompt-based LLMs via reference score or combined injection offered a potential balance, improving reliability while retaining human-readable rationales that could support SME review.

Subgroup analyses underscored the necessity of evaluating scoring fairness alongside aggregate metrics. All evaluated AES approaches displayed varying degrees of subgroup inconsistency, reflecting a prevalent pattern in recent literature (Huang & Wilson, 2025; Huang et al., 2026; Yang et al., 2024). Notably, certain disparities observed in this study diverged in direction from prior research, even when evaluating identical approaches and demographic groups, such as prompt-based LLM performance across ELL and non-ELL subgroups (Huang & Wilson, 2025). This

Dataset	Subgroup	Metric	ML	Transformer	LLM	Ensemble
Operational A	Female	QWK	0.66	0.69	-	-
		SMD	0.11	0.18	-	-
		EA	57.10	59.86	-	-
	Male	QWK	0.66	0.71	-	-
		SMD	0.08	0.15	-	-
		EA	52.49	55.69	-	-
	ELL	QWK	0.67	0.71	-	-
		SMD	0.13	0.16	-	-
		EA	51.78	54.44	-	-
	Non-ELL	QWK	0.62	0.65	-	-
		SMD	0.06	0.18	-	-
		EA	57.83	61.12	-	-
Operational B	Female	QWK	-	-	0.42	-
		SMD	-	-	0.21	-
		EA	-	-	48.28	-
	Male	QWK	-	-	0.48	-
		SMD	-	-	0.27	-
		EA	-	-	49.85	-
	ELL	QWK	-	-	0.53	-
		SMD	-	-	0.29	-
		EA	-	-	33.33	-
	Non-ELL	QWK	-	-	0.47	-
		SMD	-	-	0.22	-
		EA	-	-	49.74	-

Table 4: Model subgroup performance on Conventions.

divergence indicates that the magnitude and direction of potential bias depend not only on the AES architecture and target subgroups, but also on specific assessment contexts and item prompts. Interpreting these patterns alongside qualitative audits of model reasoning and human scoring behaviors will be essential for identifying underlying causes and promoting equitable scoring.

Beyond predictive performance, operational demands varied considerably across approaches. ML and transformer models require substantial labeled training data and full model retraining for updates. In contrast, prompt-based LLMs facilitate rapid, low-resource iterations via prompt optimization. Ultimately, the transparent rationales offered by LLM and ensemble architectures provide key advantages for system validation, auditing, and stakeholder trust.

A primary limitation of this study is that LLM prompt optimization was conducted exclusively on the TD scale. The CV scale remained unoptimized, leading to incomplete prompt-based LLM and ensemble results that complicate comparisons against other approaches. Future research should apply consistent optimization protocols across all scoring scales to evaluate domain- and scale-specific generalizability. Additionally, while subgroup analyses identified performance discrepancies across demographics, the root causes remain unexamined. Future studies should integrate model error analyses, human scoring patterns, and student writing characteristics to isolate contributing factors and establish frameworks for mitigating subgroup variance.

6 Conclusions

This study conducted a comparative evaluation of four AES approaches, including traditional ML, transformer fine-tuning, prompt-based LLMs, and a novel ensemble approach, within a high-stakes, large-scale assessment context.

The results confirm that fine-tuned transformers deliver the highest scoring reliability and subgroup consistency, followed by traditional ML models. Meanwhile, the novel ensemble architecture bridges the gap between ML reliability and LLM explainability, demonstrating that hybrid

architectural designs can effectively leverage the complementary strengths of both approaches.

Finally, operational readiness in AES requires evaluating models beyond raw reliability to ensure fairness, auditability, transparency, and defensible decision-making.

References

Atkinson, John, & Palma, Diego. (2025). An LLM-based hybrid approach for enhanced automated essay scoring. *Scientific Reports, 15*(1). https://doi.org/10.1038/s41598-025-87862-3.

Chu, SeongYeub, Kim, JongWoo, Wong, Bryan, & Yi, MunYong. (2025). Rationale behind essay scores: Enhancing S-LLM's multi-trait essay scoring with rationale generated by LLMs. *arXiv*. https://arxiv.org/abs/2410.14202.

Hou, Zhaoyi Joey, Ciuba, Alejandro, & Li, Xiang Lorraine. (2025). Improve LLM-based automatic essay scoring with linguistic features. *arXiv*. https://arxiv.org/abs/2502.09497.

Huang, Yue, & Wilson, Joshua. (2025). Evaluating LLM-based automated essay scoring: accuracy, fairness, and validity. *Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con)* (pp. 71–83). National Council on Measurement in Education (NCME). https://aclanthology.org/2025.aimecon-wip.9.pdf.

Huang, Yue, Palermo, Corey, & Wilson, Joshua. (2026). Accuracy and fairness of generative AI in automated essay scoring: Comparing GPT-4o, feature-based models, and human raters. *Assessing Writing*, *69*. https://doi.org/10.1016/j.asw.2026.101047.

Lee, Sanwoo, Cai, Yida, Meng, Desong, Wang, Ziyang, & Wu, Yunfang. (2024). Unleashing large language models' proficiency in zero-shot essay scoring. *arXiv*. https://arxiv.org/abs/2404.04941.

Li, Shengjie, & Ng, Vincent. (2024). Automated essay scoring: a reflection on the state of the art. *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing* (pp. 17876–17888). Association for Computational Linguistics. http://doi.org/10.18653/v1/2024.emnlp-main.991.

Ramesh, Dadi, & Sanampudi, Suresh Kumar. (2022). An automated essay scoring systems: a systematic literature review. *Artificial Intelligence Review*, *55*, 2495–2527. https://doi.org/10.1007/s10462-021-10068-2.

Stahl, Maja, Biermann, Leon, Nehring, Andreas, & Wachsmuth, Henning. (2024). Exploring LLM prompting strategies for joint essay scoring and feedback generation. *arXiv*. https://arxiv.org/abs/2404.15845.

Wang, Yongjie, Wang, Chuang, Li, Ruobing, & Lin, Hui. (2022). On the use of Bert for automated essay scoring: Joint learning of multi-scale essay representation. *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (pp. 3416–3425). Association for Computational Linguistics. https://doi.org/10.18653/v1/2022.naacl-main.249.

Xiao, Changrong, Ma, Wenxing, Song, Qingping, Xu, Sean Xin, Zhang, Kunpeng, Wang, Yufang, & Fu, Qi. (2024). Human-AI collaborative essay scoring: A dual-process framework with LLMs. *arXiv*. https://arxiv.org/abs/2401.06431.

Yancey, Kevin P., Laflair, Geoffrey, Verardi, Anthony, & Burstein, Jill. (2023). Rating short L2 essays on the CEFR scale with GPT-4. *Proceedings of the 18th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2023)* (pp. 576–584). Association for Computational Linguistics. https://doi.org/10.18653/v1/2023.bea-1.49.

Yang, Kaixun, Raković, Mladen, Li, Yuyang, Guan, Quanlong, Gašević, Dragan, & Chen, Guangliang. (2024). Unveiling the tapestry of automated essay scoring: A comprehensive investigation of accuracy, fairness, and generalizability. *Proceedings of the AAAI Conference on Artificial Intelligence*, *38*(20), 22466–22474. https://doi.org/10.1609/aaai.v38i20.30254.

Yang, Ruosong, Cao, Jiannong, Wen, Zhiyuan, Wu, Youzheng, & He, Xiaodong. (2020). Enhancing automated essay scoring performance via fine-tuning pre-trained language models with combination of regression and ranking. *Findings of the Association for Computational Linguistics: EMNLP 2020* (pp. 1560–1569). Association for Computational Linguistics. https://doi.org/10.18653/v1/2020.findings-emnlp.141.

Zhang, Su, Hove, Jennifer, & Gu, Zhimei. (2026, October 5-7). *A jurisdictional scan of automated scoring practices in K-12 summative assessment programs* [Conference presentation]. Artificial Intelligence in Measurement and Education Conference, Pittsburgh, PA, United States.

Construct Validity of Small-Sample Transformer Scoring Models: A Mechanistic Interpretability Approach

Michael P. Hemenway and Martha Bellows
Pearson
{michael.hemenway, martha.bellows}@pearson.com

Abstract

Transformer scoring models fine-tuned on small samples (n=64) can reach deployment-quality agreement with human raters, but agreement statistics cannot reveal whether models lean harder on construct-irrelevant surface features like response length. This paper explores Mechanistic Interpretability (MI) strategies as a framework for auditing scoring engines, evaluating six diagnostic techniques across 90 DeBERTa-v3-base models (6 ASAP2 prompts $\times$ 3 training sizes $\times$ 5 seeds). We distinguish between *correlational* techniques (saliency maps, probing, activation summaries) and *interventional* techniques (counterfactual edits, representational erasure). Correlational methods show suggestive trends—small-sample models exhibit higher length partial correlations ($d = 0.65$) and surface-dominated layer representations ($d = 0.60$)—but suffer from seed instability (80–96% seed noise) and lack power under multiplicity corrections ($q = .17$). Interventional methods provide definitive proof: surgically deleting length directions from a model's 768-dimensional internal representation destroys 88% of agreement at n=64 versus 67% at n=1024 ($d = -0.75$, $p = .005$, $q = .049$), an effect independent of accuracy ($r = -0.018$). Counterfactual padding shows that duplicating sentences raises predicted scores for 86–92% of responses across all sizes. We synthesize these findings into an actionable MI audit protocol for measurement practitioners.

1 Introduction

Automated scoring engines are increasingly built by fine-tuning pre-trained transformer language models on small sets of human-scored responses. Recent work demonstrates that with systematic hyperparameter optimization, deployment-quality agreement (quadratic weighted kappa, QWK $\geq$ 0.70) is attainable from as few as 64 training responses (Barber et al., 2025). The operational appeal is obvious: human scoring is the dominant cost in constructed-response assessment.

However, high agreement with human raters does not guarantee construct validity. Agreement is a purely behavioral criterion, and a model can reproduce human scores by exploiting surface heuristics that no scoring rubric endorses. Response length is the classic example: because length correlates strongly with human essay scores ($r \approx 0.70$ pooled), a length-only scorer can achieve non-trivial agreement (Perelman, 2014), and adversarial manipulations can game automated engines (Powers et al., 2002). When fine-tuned on small samples, models have fewer examples to discover complex construct signals and stronger incentives to rely on superficial shortcuts. Standard agreement metrics cannot detect when this occurs.

Introducing Mechanistic Interpretability to Measurement. In an attempt to move from "black box" to "glass box" models, the machine learning discipline has developed *Mechanistic Interpretability* (MI)—a subfield focused on reverse-engineering the internal representations, circuits, and decision logic of neural networks.

This paper serves a dual purpose for measurement scientists and educational ML practitioners. Substantively, we answer a pressing operational question: *Do small-sample scoring models rely more on construct-irrelevant surface features than models trained on larger samples?* Methodologically, we evaluate a battery of MI techniques, illustrating which strategies yield actionable validity evidence and which fail due to noise, unanchored activations, or uncalibrated baselines.

Scope and Validity Argument. We do not claim that MI replaces a comprehensive validity argument. Validity is a property of score interpretation and use, supported by evidence regarding

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 322–328
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

test content, response processes, internal structure, and consequences (Messick, 1989; Kane, 2013; American Educational Research Association et al., 2014). MI specifically informs two strands: *response processes* (how input features influence scores) and *internal structure* (how concepts are encoded in model representations).

We evaluate 90 models spanning six prompts, three training sizes ($n \in \{64, 512, 1024\}$), and five random seeds, analyzing them with six MI approaches categorized into *correlational* and *interventional* strategies.

2 Mechanistic Interpretability Strategies

We group MI techniques into three categories based on what they measure and why a measurement scientist might use them.

2.1 Data and Model Setup

We analyze the ASAP2 corpus of human-scored essays (scale 1–6) across six diverse prompts (Crossley et al., 2025). For each prompt, we fine-tune `microsoft/deberta-v3-base` (He et al., 2023) on training splits of $n \in \{64, 512, 1024\}$ across 5 independent random seeds ($6 \times 3 \times 5 = 90$ models). Validation (n=400) and test sets (n=200) are held fixed per prompt. Models are optimized using Optuna TPE search over 30 trials (Akiba et al., 2019) with frequency-balanced MSE loss, and mapped to score points via post-hoc affine calibration.

2.2 Category A: Feature Association and Probing

Partial Correlations (Behavioral Association). *Rationale*: We want to know if model predictions align with surface features (word count, character count, punctuation) beyond what is justified by human scores. We compute the partial correlation $r(\hat{y}, x_j \mid y)$ between model predictions $\hat{y}$ and surface feature x_j, partialling out the human score y. Non-zero partial correlations indicate prediction variance tied to surface shortcuts that human raters do not endorse.

Probing Classifiers (Representational Content). *Rationale*: Probing asks what information is *stored* inside intermediate neural layers (Belinkov, 2022). We extract mean-pooled hidden activations at layers $\{0, 3, 6, 9, 12\}$ and train logistic regression probes to predict quintile buckets of surface features versus construct targets (human score quintiles). The *probing gap* (surface accuracy minus construct accuracy) tests whether internal representations carry more information about response appearance than essay quality.

2.3 Category B: Saliency and Activation Summaries

Token Attributions (Integrated Gradients & ADI). *Rationale*: Attribution methods assign each token a score reflecting its contribution to the final prediction (Sundararajan et al., 2017), computed here using Captum (Kokhlikyan et al., 2020). Measurement researchers often view attribution maps as automated rubric checks. We compute the *Surface Feature Reliance Score* (SFRS) as the share of attribution mass landing on non-content tokens (function words and punctuation). To test consistency, the *Attribution Divergence Index* (ADI) measures token-attribution dissimilarity ($1 - \rho_{\text{Spearman}}$) between two models scoring the exact same response.

Attention Patterns and Logit Lens. *Rationale*: Attention heads indicate where models focus, while the *logit lens* (Nostalgebraist, 2020) reads out intermediate predictions at each layer to see how early a score settles. We track attention entropy, content-token focus, and the layer at which predictions reach stability.

2.4 Category C: Interventional Causal Methods

Counterfactual Length Manipulation. *Rationale*: Correlational methods show association, not causal use. To test if models *change scores* when length changes, we create sentence-aligned inputs: **Base** (original prefix), **Short** (randomly deletes 30% of sentences), and **Restored** (duplicates remaining sentences to match Base's length without adding information). The *Length Compensation Effect* (LCE) is the score difference between Restored and Short: score points awarded for length alone.

Representational Erasure (Amnesic Probing). *Rationale*: Before predicting a score, DeBERTa compresses a response into a 768-dimensional vector. Using iterative nullspace projection (Ravfogel et al., 2020; Elazar et al., 2021), we identify the directions in this 768-dimensional space encoding response length and surgically remove them before passing the edited vector to the scoring layer. Comparing agreement loss against a

	QWK_{raw}	QWK_{cal}	Exact	≥ 0.70
n=64	0.687	0.718 (0.056)	0.561	20/30
n=512	0.743	0.771 (0.036)	0.593	29/30
n=1024	0.759	0.787 (0.030)	0.602	30/30

Table 1: Test-set agreement across training sizes (mean over 30 models, SD in parentheses). Small-sample models frequently reach deployment thresholds (QWK ≥ 0.70), setting up our interpretability comparison.

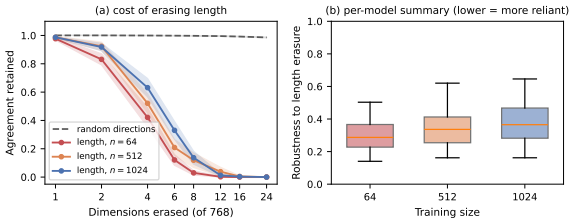


Figure 1: Representational erasure. (a) Agreement retained as length-encoding directions are removed from the 768-dimensional vector, compared to erasing random directions. (b) Per-model robustness ratio (retained agreement under length erasure vs. random control).

control that erases random directions tests whether length representation is causally load-bearing.

3 Substantive Findings

3.1 Behavioral Performance Baseline

Table 1 establishes the baseline. Mean calibrated QWK rises from 0.718 at n=64 to 0.787 at n=1024, with 20 of 30 small-sample models clearing the 0.70 deployment threshold. Because accuracy differs across sample sizes, we control for QWK throughout our analyses to ensure surface reliance is not merely a proxy for lower agreement.

3.2 Causal Evidence: Length is Load-Bearing

The most decisive finding comes from representational erasure (Figure 1 and Table 2). Erasing random directions from the 768-dimensional pooled vector is nearly harmless, retaining 98–99% of QWK at all sample sizes. In contrast, erasing length directions causes a catastrophic drop in agreement—and differentially so. Erasing just six length directions destroys 88% of human agreement for n=64 models (12% retained) versus 67% for n=1024 models (33% retained).

Across training sizes, the overall robustness

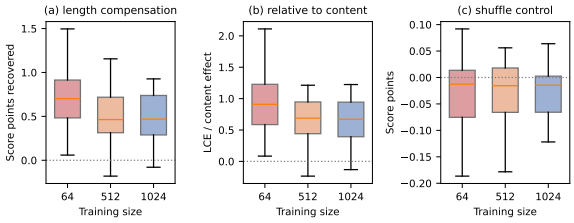


Figure 2: Counterfactual length manipulation. (a) Score points recovered by duplicating sentences. (b) Share of content-deletion penalty recovered. (c) Sentence-shuffle control.

score (retained agreement under length erasure divided by random control) rises from 0.299 at n=64 to 0.379 at n=1024 ($d = -0.75, p = .005, q = .049$). Crucially, this robustness score is uncorrelated with test QWK ($r = -0.018$). When restricting analysis to an accuracy-matched band ($0.697 \leq$ QWK ≤ 0.804), small-sample models remain significantly more length-dependent ($d = -0.81, p = .003$). This proves that small-sample models causally depend far more on internal length representations, independent of model accuracy.

3.3 Causal Evidence: Vulnerability to Surface Padding

The counterfactual edit provides a direct estimate of gaming vulnerability in score points (Figure 2 and Table 2). Restoring deleted length with duplicated sentences recovers +0.68 score points at n=64 compared to +0.48 points at n=1024 ($d = +0.54, p = .042$). Across all training sizes, 86–92% of individual essay predictions increase when padded with restatements, returning 89% of the content-deletion penalty at n=64 against 63% at n=1024.

Two controls confirm specificity: deleting content causes a uniform score drop (0.75–0.78 points across sizes), and shuffling sentence order has virtually no impact (-0.04 points), proving models respond to length and redundancy rather than general perturbation.

3.4 Correlational Evidence: Association and Probing

The correlational methods align with the causal findings but carry weaker statistical power.

Partial Correlations. Residual length correlation (Figure 3) drops from 0.555 at n=64 to 0.460 at n=1024 ($d = 0.65, p = .014$). Matching mod-

Measure	n=64	n=1024	d	90% CI	p_{diff}	q_{BH}	Verdict
Length compensation effect (score pts)	0.680	0.484	+0.54	[+0.11, +0.97]	.042	.087	suggestive
Length compensation per 100 tokens	0.545	0.393	+0.52	[+0.09, +0.95]	.048	.087	suggestive
Length reliance ratio	0.887	0.630	+0.52	[+0.09, +0.95]	.048	.087	suggestive
Content effect (score pts)	0.775	0.752	+0.18	[-0.26, +0.61]	.497	.527	inconclusive
Sentence-shuffle effect (score pts)	-0.031	-0.044	+0.16	[-0.27, +0.60]	.527	.527	inconclusive
Directions lost before agreement halves	4.933	6.133	-0.59	[-1.02, -0.16]	.025	.087	suggestive
Robustness to length erasure	0.299	0.379	-0.75	[-1.18, -0.31]	.005	.049	difference
Agreement lost after erasing 8 directions	0.704	0.677	+0.18	[-0.26, +0.61]	.497	.527	inconclusive
r(prediction, word count)	0.777	0.753	+0.36	[-0.07, +0.79]	.170	.256	inconclusive

Table 2: Causal interventional experiments comparing n=64 and n=1024 models. Rows 1–5 report counterfactual length manipulations; rows 6–9 report amnesic representational erasure. Positive d indicates greater length reliance at n=64 (except rows 6–7 where lower values indicate greater reliance).

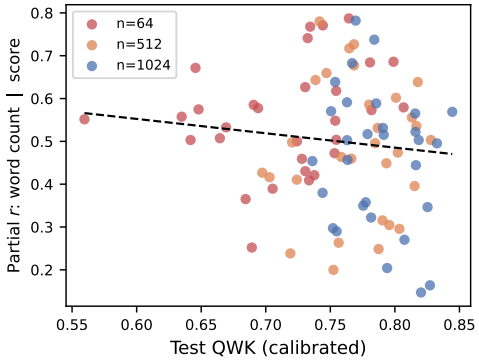


Figure 3: Residual length reliance (partial correlation between prediction and word count, controlling for human score) versus test QWK across all 90 models.

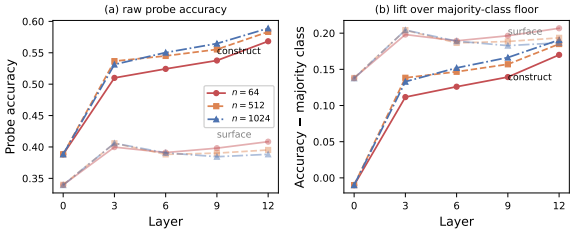


Figure 4: Probing classifier accuracy by layer. (a) Raw accuracy. (b) Lift over majority-class chance floor (construct floor = 0.398, surface floor = 0.202).

els on QWK preserves this difference (0.577 vs 0.481, $d = 0.64, p = .022$), confirming that two models with identical human agreement can differ in surface reliance.

Probing Gap. Scoring layer probes as lift over chance floors (Figure 4) shows that the probing gap shrinks from +0.148 at layer 0 (frozen embeddings) to +0.037 at layer 12 for n=64, +0.009 for n=512, and −0.004 for n=1024 ($d = 0.60, p = .023$). Fine-tuning converts a surface-dominated representation into a construct-dominated one, but only larger training sets complete this conversion at the output layer.

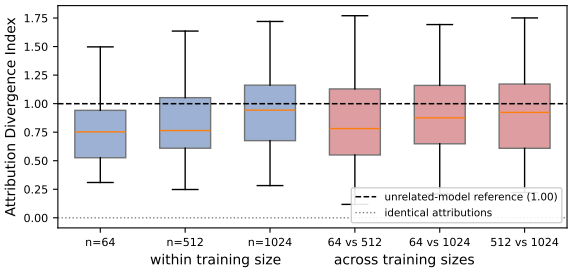


Figure 5: Attribution Divergence Index (ADI) between model seeds vs. between training conditions. Token attributions between seeds are as divergent (0.853) as attributions between sample sizes (0.881).

4 Methodological Audit of MI Strategies

A key goal of this study is evaluating MI tools for educational assessment settings. Table 3 synthesizes our findings into a diagnostic guide for automated scoring contexts.

Lesson 1: Token Saliency Maps Suffer from Seed Instability. Gradient-based token attributions (Integrated Gradients) are often promoted as automated rater checks. However, comparing attribution vectors across models reveals extreme instability (Figure 5). The divergence between two models trained on identical data with different random seeds (ADI = 0.853) is virtually indistinguishable from the divergence between small and large models (ADI = 0.881, $d = 0.08, p = .350$). Saliency maps reflect random initialization noise as much as learned representations (Kindermans et al., 2019; Adebayo et al., 2018). *Recommendation*: Token attribution maps must never be interpreted without reporting a same-condition seed baseline.

MI Strategy	Measurement Rationale	Empirical Result (n=64 vs 1024)	Diagnostic Utility & Recommendations
Partial Correlations *(Residual Alignment)*	Test if predictions align with surface features after controlling for human scores.	Monotonic decrease: $r = .555 \rightarrow .460$ ($d = 0.65, p = .014$).	**Suggestive clue**: Measures association, not causation; vulnerable to length-construct collinearity and battery multiplicity ($q = .17$).
Probing Classifiers *(Layer Representation)*	Audit whether internal layers encode surface features vs. construct targets.	Chance-corrected gap: $+.078 \rightarrow +.054$ ($d = 0.60, p = .023$).	**Requires chance floors**: Raw accuracy is misleading due to score distribution skew. Score probes as lift over majority-class floor.
Token Attributions *(Saliency / ADI / SFRS)*	Inspect token-level attributions for function/punctuation vs content tokens.	Seed vs seed ADI $= 0.853$; condition ADI $= 0.881$ ($d = 0.08$).	**High seed instability**: Saliency maps vary wildly across random seeds (80–96% seed noise). Uninformative without a seed baseline.
Attention & Logit Lens *(Internal Flow)*	Track head entropy, content focus, and intermediate prediction settlement.	Content mass $d = -0.08$; stabilization layer $d = +0.09$.	**Unanchored activations**: Activation summaries lack task controls; power was inadequate for unconstrained summaries ($d_{\text{min}} = 0.74$).
Counterfactual Edits *(Length Padding)*	Test if adding length without information alters predicted scores.	Score recovery: $+0.68$ vs $+0.48$ pts; 86–92% gain score from padding.	**Actionable in score points**: Directly measures model gaming vulnerability and susceptibility to restatements.
Representational Erasure *(Amnesic Probing)*	Surgically delete length directions from 768-d pooled representation.	Agreement retained: 12% vs 33% ($d = -0.75, p = .005, q = .049$).	**Causally diagnostic**: Proves length is load-bearing; independent of QWK ($r = -0.018$) and holds under QWK-matching ($d = -0.81$).

Table 3: Diagnostic evaluation of Mechanistic Interpretability (MI) techniques for measurement applications. Interventional strategies provide actionable validity evidence, whereas unanchored correlational summaries suffer from seed noise and multiplicity.

Lesson 2: Probing Classifiers Require Chance Floors. Raw probing accuracy suggested models encode construct better than surface features. However, construct targets (score quintiles collapsed to 4 categories) carried an artificial chance floor of 0.398, whereas surface targets (balanced quintiles) carried a floor of 0.202. Raw probing accuracy was completely uninterpretable. Subtracting chance floors reversed the conclusion, showing surface dominance in early and middle layers (Figure 4b), aligning with control-task principles for probing (Hewitt and Liang, 2019; Voita and Titov, 2020). *Recommendation*: Probing accuracies must always be reported as lift over majority-class chance baselines.

Lesson 3: Unconstrained Batteries Suffer Multiplicity. Running a broad battery of 15 correlational metrics across activations without a pre-specified primary endpoint leads to severe multiplicity penalties under procedures such as Benjamini–Hochberg (Benjamini and Hochberg, 1995). While individual partial correlations and probing gaps showed nominal significance ($p \leq .023$), none survived false-discovery-rate correction ($q = .17$). In contrast, causal interventions yield strong, unconfounded signals that survive correction.

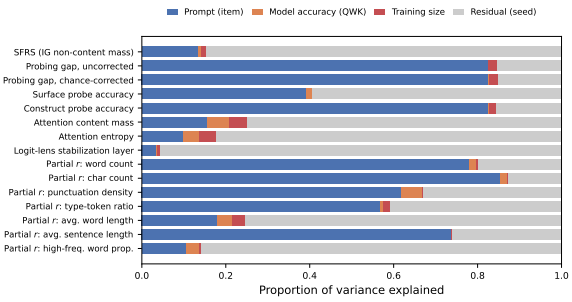


Figure 6: Variance decomposition across prompt identity, model accuracy (test QWK), training size, and seed noise.

Lesson 4: Item Identity Dominates Variance. Partitioning variance across all metrics (Figure 6) reveals that prompt/item identity accounts for 50–83% of total variance, whereas training size accounts for only 0.1–5%. *Recommendation*: Audit protocols must evaluate models across diverse items rather than pooling across prompts.

5 Practical Guidance for Practitioners

Actionable Audit Protocol. For measurement programs evaluating automated scoring engines, we recommend a 4-step validity audit:

1. **Measure Seed Noise**: Establish same-condition seed variance before drawing conclusions about model differences.

2. **Apply Counterfactual Padding Tests**: Test model predictions against restatements and duplicated text to quantify gaming vulnerability in score points.

3. **Perform Representational Erasure**: Use nullspace projection on pooled embeddings to verify whether surface shortcuts are causally load-bearing, independent of QWK.

4. **Evaluate Item-Level Variation**: Organize validity checks by item rather than relying on aggregate summary metrics.

Limitations. ASAP2 carries a single human score per response, preventing human–human baseline comparisons. While ASAP2 includes student demographic metadata (e.g., race/ethnicity, ELL status, disability, and economic status), subgroup fairness analysis was beyond the scope of this mechanistic length-reliance audit; future work will conduct intersectional subgroup analyses. Erasure removes linear directions, serving as a lower bound on feature dependence.

6 Conclusion

Mechanistic interpretability provides powerful tools for validating automated scoring engines, but implementation must be carefully designed and can benefit from collaboration between AI Scientists and Psychometricians. Small-sample transformer scoring models (n=64) reach high human agreement but causally depend more on construct-irrelevant surface length representations than models trained on larger samples. While passive correlational summaries suffer from seed noise and multiplicity, interventional techniques—specifically representational erasure and counterfactual edits—offer robust, actionable validity evidence for assessment.

References

Julius Adebayo, Justin Gilmer, Michael Muelly, Ian Goodfellow, Moritz Hardt, and Been Kim. 2018. Sanity checks for saliency maps. In *Advances in Neural Information Processing Systems 31 (NeurIPS 2018)*, pages 9505–9515, Montréal, Canada.

Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. 2019. Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 2623–2631.

American Educational Research Association, American Psychological Association, and National Council on Measurement in Education. 2014. *Standards for Educational and Psychological Testing*. American Educational Research Association, Washington, DC.

Justin O. Barber, Michael P. Hemenway, and Edward Wolfe. 2025. When does active learning actually help? Empirical insights with transformer-based automated scoring. In *Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con): Coordinated Session Papers*, pages 1–8, Pittsburgh, Pennsylvania. National Council on Measurement in Education.

Yonatan Belinkov. 2022. Probing classifiers: Promises, shortcomings, and advances. *Computational Linguistics*, 48(1):207–219.

Yoav Benjamini and Yosef Hochberg. 1995. Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B (Methodological)*, 57(1):289–300.

Scott A. Crossley, Perpetual Baffour, L. Burleigh, and Jules King. 2025. A large-scale corpus for assessing source-based writing quality: ASAP 2.0. *Assessing Writing*, 65:100954.

Yanai Elazar, Shauli Ravfogel, Alon Jacovi, and Yoav Goldberg. 2021. Amnesic probing: Behavioral explanation with amnesic counterfactuals. *Transactions of the Association for Computational Linguistics*, 9:160–175.

Pengcheng He, Jianfeng Gao, and Weizhu Chen. 2023. Debertav3: Improving deberta using electra-style pre-training with gradient-disentangled embedding sharing. *Preprint*, arXiv:2111.09543.

John Hewitt and Percy Liang. 2019. Designing and interpreting probes with control tasks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2733–2743, Hong Kong, China. Association for Computational Linguistics.

Michael T. Kane. 2013. Validating the interpretations and uses of test scores. *Journal of Educational Measurement*, 50(1):1–73.

Pieter-Jan Kindermans, Sara Hooker, Julius Adebayo, Maximilian Alber, Kristof T. Schütt, Sven Dähne, Dumitru Erhan, and Been Kim. 2019. The (un)reliability of saliency methods. In Wojciech Samek, Grégoire Montavon, Andrea Vedaldi, Lars Kai Hansen, and Klaus-Robert Müller, editors, *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*, volume 11700 of *Lecture Notes in Computer Science*, pages 267–280. Springer, Cham.

Narine Kokhlikyan, Vivek Miglani, Miguel Martin, Edward Wang, Bilal Alsallakh, Jonathan Reynolds, Alexander Melnikov, Natalia Kliushkina, Carlos Araya, Siqi Yan, and Orion Reblitz-Richardson. 2020. Captum: A unified and generic model interpretability library for PyTorch. arXiv preprint arXiv:2009.07896.

Samuel Messick. 1989. Validity. In Robert L. Linn, editor, *Educational Measurement*, 3rd edition, pages 13–103. American Council on Education / Macmillan, New York, NY.

Nostalgebraist. 2020. Interpreting GPT: The logit lens. LessWrong.

Les Perelman. 2014. When "the state of the art" is counting words. *Assessing Writing*, 21:104–111.

Donald E. Powers, Jill C. Burstein, Martin Chodorow, Mary E. Fowles, and Karen Kukich. 2002. Stumping e-rater:challenging the validity of automated essay scoring. *Computers in Human Behavior*, 18(2):103–134.

Shauli Ravfogel, Yanai Elazar, Hila Gonen, Michael Twiton, and Yoav Goldberg. 2020. Null it out: Guarding protected attributes by iterative nullspace projection. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7237–7256, Online. Association for Computational Linguistics.

Mukund Sundararajan, Ankur Taly, and Qiqi Yan. 2017. Axiomatic attribution for deep networks. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, ICML'17, pages 3319–3328. JMLR.org.

Elena Voita and Ivan Titov. 2020. Information-theoretic probing with minimum description length. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 183–196, Online. Association for Computational Linguistics.

Applying Evidence-Centered Design to Automated Evals of AI-Powered Assessment Systems

Kristen DiCerbo
Khan Academy
kdicerbo@cox.net

Britte Haugan Cheng
Menlo Education Research
bcheng@menloedu.org

John Whitmer
Learning Data Insights
john@ld-insights.com

Abstract

This paper demonstrates an application of Evidence-Centered Design (ECD) as a principled approach to design automated evaluations of AI-powered assessment outputs. We demonstrate this application through Khan Academy's "Explain Your Thinking" conversational agent for mathematics items, showing how ECD's layered models can be translated into rigorous and interpretable systems to demonstrate validity, reliability and fairness in AI systems to diverse audiences.

1 Introduction

The emergence of generative AI in educational assessment has expanded what is possible: real-time conversational probing of student reasoning, automated criterion-by-criterion scoring, and dynamic supervision of AI-generated instructional dialogue. These capabilities bring genuine potential — and a genuine validation problem. As AI systems take on roles previously reserved for human scorers and interviewers, generating follow-up questions, interpreting student responses, and assigning scores in real time, the field needs rigorous frameworks for designing and validating the performance of these systems.

Without such frameworks, automated evaluations of generative AI outputs, commonly called "evals", risk becoming ad hoc checklists that test for surface-level accuracy without grounding claims in a defensible theory of evidence. This is not merely a theoretical concern. A growing body of work on automated scoring and LLM-based evaluation has demonstrated that agreement with human raters, while necessary, is insufficient as a standalone validity criterion: what counts as adequate evidence, and for what claims, must be specified before evaluation begins (Kane, 2006; Mislevy et al., 2003b).

Evidence-Centered Design (ECD), developed by Mislevy and colleagues, was designed precisely to build validity arguments for claims (Mislevy et al., 2003a,b). ECD frames assessment as "an exercise in evidentiary reasoning." Its layered architecture explicitly connects the claims being made about a student, the evidence required to support those claims, and the task conditions under which that evidence is gathered. Crucially, ECD builds validity into design from the outset rather than treating it as a post-hoc audit.

These properties make ECD a natural scaffold not only for designing assessments, but also for designing the evaluations that determine whether AI components within those assessments are functioning as intended. Yet despite ECD's established role in assessment design, its application as a framework specifically for AI component evaluation, for structuring the claims, evidence, and tasks involved in evaluating AI scorer and conversationalist components, has not been systematically explored. This paper addresses that gap.

2 Background

2.1 Evidence-Centered Design

ECD is a comprehensive framework for assessment design introduced by Mislevy, Steinberg, and Almond (2003). Its conceptual roots draw on developments in Bayesian inference, cognitive science, and legal evidentiary reasoning, including Wigmore's and Toulmin's frameworks for structuring arguments from evidence to claim (Mislevy et al., 2003b; Toulmin, 1958).

The core of ECD's Conceptual Assessment Framework (CAF) consists of four interrelated models:

- The Student Model specifies what claims are being made about the examinee — the knowledge, skills, and abilities (KSAs) the assessment is designed to measure. These are expressed as variables that can take on different

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 329–334
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

values and that can, in principle, be updated by evidence.

- The Evidence Model connects observations to claims. It has two sub-components: the evaluation component, which describes what features of a student's work product are relevant to each KSA, and the measurement model, which specifies how those features are translated into updated beliefs about the student model variables.

- The Task Model specifies the conditions under which evidence is elicited: the characteristic features of tasks that are relevant to the targeted KSAs, and the variable features that can be manipulated.

By requiring explicit articulation of each model and the connections between them, it forces designers to specify the full inferential chain from task performance to score-based claim, making the reasoning transparent, auditable, and extensible/improvable (Mislevy and Riconscente, 2006; Mislevy et al., 2012).

2.2 The Problem of AI Eval Design

Generative AI is increasingly being used to create new assessment items and experiences. The proliferation of AI components in educational assessment has created a new class of evaluation problems. An AI scoring application is not a test, it is a prediction of a score that must itself be evaluated. An AI conversation agent is not a student, it is an algorithm that predicts likely responses based on observed patterns whose output quality must be assessed against specified criteria. Compounding this factor, generative AI has built-in variation as a feature to create more dynamic, human-sounding language. The field has not yet developed principled, reusable frameworks for this class of evaluation problem.

Current practice in AI eval design tends toward one of two options. First, benchmark-based evaluation assess AI components against standardized datasets without specifying the inferential chain linking benchmark performance to real-world claims. Most benchmarks are built around predicting test responses, not identifying learning-related behaviors nor competencies. uses. Second, ad hoc validation creates bespoke evaluation datasets and metrics, but often without a principled framework to ensure that the evidence gathered is sufficient for the claims being made. Too often, these evaluations choose what is easy to measure using existing algorithmic techniques, instead of selecting the most important behavior to measure to demonstrate a claim.

Both approaches lack the explicit links between claims and evidence. ECD's framework addresses this antecedent question directly.

3 Context: Khan Academy's Explain Your Thinking System

Khan Academy's Explain Your Thinking (EYT) feature provides the empirical context for this paper. EYT engages grade 3–12 students in multi-turn AI-guided conversations about their mathematical reasoning after they solve a problem. The system's purpose is to elicit and assess the conceptual understanding underlying a student's correct or incorrect answer — a diagnostic function that goes beyond what a single response can reveal. Currently used in high school Algebra and Geometry with approximately 10,000 students, EYT extends well-established research on conversation-based assessment (Jackson and Zapata-Rivera, 2016; Liu et al., 2016) by using generative AI to operationalize conversations at scale — replacing the brittle, pre-scripted dialogue trees that constrained earlier systems with flexible prompt-based control (DiCerbo et al., 2026).

3.1 System Architecture

EYT employs a modular architecture with three distinct AI components:

- The AI Scorer evaluates each student response against a set of analytic, dichotomously-scored criteria authored by content experts. Each criterion represents a specific aspect of conceptual understanding targeted by the item. The Scorer's output, a binary judgment for each criterion, drives the conversation.

- The Conversation Agent generates targeted follow-up questions based on the criteria that the Scorer has identified as unmet. Its role is to probe the student's reasoning in a way that mirrors what a skilled teacher might do: asking questions that are targeted, non-leading, and appropriate to the student's demonstrated level of understanding.

- The Supervisor Agent reviews each question generated by the Conversation Agent before it

is presented to the student. Its role is to ensure that questions do not inadvertently provide the answer, violate content policies, or otherwise compromise the integrity of the assessment.

Together, these three components implement a complete assessment cycle: scoring, targeted questioning, and quality assurance. Each component makes consequential decisions that depend on the outputs of the others, creating a tightly coupled system in which validation failures in one component propagate through the whole.

3.2 EYT Validation Practice

Current validation practice for EYT has involved constructing synthetic student response datasets, establishing human-labeled ground truth through expert annotation, and iterating on scoring rubrics through pre-pilot validation cycles before advancing to pilot testing with authentic student responses collected from live use with approximately 10,000 high school mathematics students.

To make these practices concrete, a human content creator writes an item asking students to select which triangle is similar to a given triangle. They then write criteria for the EYT conversation about similarity, including, "Student understands that the two triangles are similar by AA similarity criteria (i.e., two of the angle pairs are the same). Similarity must be stated and not merely implied, and the angles must also be mentioned as a basis for similarity." Following this, synthetic conversations responding to the item are developed. Humans then label these conversations as meeting or not meeting the criteria. In places where there is significant human disagreement, the criteria are revised and the process started again. When there is human agreement, AI scoring prompts are developed and iterated on until they reach sufficient reliability.

Human agreement on criterion judgments is measured using Matthews Correlation Coefficient (MCC) and Cohen's kappa to establish a reliability ceiling before AI agreement is assessed. This sequencing, establishing human-human reliability before measuring human-AI agreement, reflects an important methodological principle that ECD makes explicit.

4 Methods: Mapping ECD to AI Eval Design

We conduct a structured conceptual mapping between ECD's framework and the practice of AI eval design, grounded in the EYT case. For each ECD model, we identify the corresponding eval design question and demonstrate how EYT's validation approach addresses each question, or would benefit from explicitly addressing it. The mapping is summarized in Table 1 (below).

4.1 The Student Model

In an AI eval context, the Student Model becomes what we term the AI Component Model: an explicit specification of the claims being made about the AI component's behavior. This is the first discipline ECD imposes, and it is frequently skipped in practice.

For EYT's AI Scorer, the claim is that its binary criterion judgments correspond to those an expert human rater would make when applying the same rubric to the same response. This claim has an important internal structure: it is made per criterion, not globally. A scorer that performs well on high-frequency, clear-cut criteria but poorly on low-frequency, ambiguous ones has a very different validity profile than one that performs uniformly, a distinction that aggregate accuracy statistics obscure. ECD's student model architecture, which represents claims as a structured set of variables rather than a single score, provides the representational vocabulary to make this distinction explicit.

There are currently two claims for the Conversation Agent: (a) targeted at one or more unmet criteria, (b) not leading or outright giving the answer. Each of these sub-claims requires its own evidence model. In the future, other criteria may be explored, such as aligning to the student's demonstrated level of reasoning, requiring the development of additional evidence models, or other items..

4.2 The Evidence Model

ECD's evidence model has two sub-components that map directly onto the two central decisions in eval dataset construction: what to label, and how to aggregate those labels into a measurement.

The evaluation component maps onto the design of the labeling rubric and the annotation task. In EYT's pre-pilot validation, human expert labels serve as Observable Variables in ECD's terminology: observable expressions of the targeted construct (criterion mastery) that provide the evidentiary basis for evaluating the AI component's claims. The quality of these labels and therefore the quality of the entire evaluation depends on the quality of the annotation rubric and the expertise

ECD Model	AI Eval Analogue	Core Design Question	EYT Instantiation
Student Model	AI Component Model	What claims are being made about the component's behavior?	Scorer: binary criterion judgments match expert human rater; Conversation Agent: questions are aligned to criteria and non-leading
Evidence Model (Evaluation Component)	Labeled Dataset Design	What observable evidence bears on which claims?	Human-labeled criterion judgments as ground truth; rubric-governed annotation
Evidence Model (Measurement Model)	Metric Selection & Sequencing	How is evidence translated into updated claims?	MCC and Cohen's kappa; human IRR established before AI agreement assessed
Task Model	Eval Dataset Design	What response conditions must be represented to support inferences?	Synthetic datasets with balanced categories, challenge sets, edge cases

Table 1: ECD Models and Their AI Eval Analogues

of the annotators. ECD's framework makes this dependency explicit and provides vocabulary for discussing it.

The measurement model maps onto the selection and interpretation of agreement statistics. ECD distinguishes between the evaluation component (the rubric that produces labels) and the measurement model (the statistical procedure that aggregates those labels into evidence about the claim). This distinction clarifies why human inter-rater reliability must be established before AI agreement is assessed: if humans applying the rubric cannot agree, the labels are not a reliable ground truth, and AI-human agreement statistics are uninterpretable. The field does not always observe this sequencing. ECD provides both the vocabulary and the logical structure to make the sequencing requirement explicit.

In EYT's context, the measurement model also requires distinguishing error types. Because false positive errors by the Scorer terminate conversation prematurely, they carry a higher cost than false negative errors, which merely prolong the conversation unnecessarily. An ECD-informed evidence model specifies this asymmetry in advance — it is a property of the measurement model, not a post-hoc discovery from pilot data.

4.3 The Task Model

ECD's task model specifies the characteristic and variable features of assessment tasks that are relevant to the targeted KSAs. In an eval context, the task model translates to the design of the evaluation dataset itself: what response types must the dataset include to support valid inference about the AI component's behavior?

EYT's use of synthetic response datasets — constructed to include balanced score categories, challenge sets targeting edge cases, and prevalence-weighted distributions — reflects implicit task modeling. ECD provides the vocabulary to make this implicit modeling explicit and replicable. In particular, ECD's distinction between characteristic features (features that make a task an instance of the relevant task type) and variable features (features whose manipulation generates evidence about specific aspects of component behavior) maps cleanly onto how datasets are constructed.. The characteristic features of the dataset include back and forth interactions between the AI tutor and student, AI tutor probing student understanding, and student responses. Variable features include the length of student response, the correctness of student responses, and completeness of individual student responses.

5 Anticipated Results

This paper reports work in progress. The full conceptual mapping, including populated examples from EYT's validation materials, will be completed in advance of the conference. Anticipated results include:

- A complete translation table mapping each

ECD construct to its AI eval analogue, populated with concrete EYT examples

- Identification of at least two cases where ECD would have prescribed a different validation sequence than was initially followed, and an assessment of the consequences
- A preliminary assessment of the generalizability of the mapping to other AI-powered assessment components beyond EYT

6 Discussion and Anticipated Contributions

6.1 ECD Improves Validity by Forcing Claim Specification

The most fundamental contribution of applying ECD to AI eval design is the discipline it imposes on claim specification. Ad hoc evals frequently conflate what is being claimed about an AI component, what evidence would support that claim, and what task conditions are required to elicit that evidence. ECD's layered architecture separates these questions and requires each to be answered before the next can be addressed.

This has direct practical consequences. An AI scorer evaluated only against a convenience sample of student responses cannot support claims about its performance on low-frequency but educationally significant response types. An AI Conversation Agent evaluated only on the fluency of its questions cannot support claims about their non-leading quality. ECD's framework makes these gaps visible before evaluation begins — not after results fail to generalize.

6.2 ECD Supports Modular Eval Design

EYT's modular component architecture — a Scorer, a Conversation Agent, and a Supervisor, each with distinct responsibilities — creates both a design opportunity and an evaluation challenge. Each component can, in principle, be evaluated independently. But the components are also coupled: the Conversation Agent's behavior is conditioned on the Scorer's output, and the Supervisor's behavior is conditioned on the Conversation Agent's output. Evaluation designs that treat the components as fully independent may miss failure modes that only emerge in the full system.

ECD's framework supports modular eval design in a way that is sensitive to this coupling. By specifying the claims for each component explicitly, and by tracing the inferential chain from each component's output to the downstream claims it supports, an ECD-informed eval design can identify the points at which component-level validity failures propagate to system-level validity failures — and prioritize evaluation effort accordingly.

6.3 ECD as a Reusable Architecture for a Field That Needs One

Beyond the EYT case, ECD offers a reusable architectural vocabulary for AI eval design that does not need to be reconstructed from first principles for each new system. The translation table in Section 4 (Table 1) is intended as a working template that developers of other AI-powered assessment systems can adapt to their own contexts.

The field of AI evaluation in education is currently characterized by proliferating evaluation frameworks and benchmarks with inconsistent theoretical grounding (cf. recent work on LLM-as-judge methodology and its reliability challenges (Schroeder and Wood-Doughty, 2025). ECD does not resolve all of these challenges, but it provides a principled starting point: a mature framework with a 20-year track record in high-stakes assessment design, whose core logic is directly applicable to the AI eval problem and whose vocabulary is already shared across the measurement science community.

6.4 Open Questions and Limitations

Several open questions remain. First, ECD was designed for assessments of human examinees; adapting its vocabulary for AI components requires conceptual work that this paper initiates but does not complete. Second, the translation from ECD's static, design-time framework to the dynamic, iterative context of AI development, where components are continuously updated and eval results feed back into model training, raises questions about how ECD's validity argument should be revised as the system evolves. Third, the relationship between ECD-informed eval design and the broader project of AI safety and alignment evaluation remains to be worked out.

Extensions to be explored in future work include: using ECD as a framework for other AI-centered design activities beyond eval; applying ECD's evidentiary reasoning structure to the automated construction and updating of evals; and examining how ECD-informed eval design interacts with emerging regulatory and legal frameworks for high-stakes AI systems.

7 Conclusion

This paper argues that ECD provides a principled, reusable architecture for designing evaluations of AI-powered assessment components, and demonstrates this argument through a structured conceptual mapping applied to Khan Academy's Explain Your Thinking system. The mapping reveals both 1) the practical value of ECD for AI eval design, forcing claim specification, supporting modular evaluation, providing reusable architectural vocabulary and 2) the conceptual work required to address the open questions that arise when a framework designed for static human assessments is applied to dynamic, interdependent AI systems.

As AI systems take on increasingly central roles in educational assessment, the field needs evaluation frameworks that are not merely technically adequate but theoretically grounded and defensible. ECD, with its roots in evidentiary reasoning and its 20-year track record in high-stakes assessment design, offers a strong foundation for that work.

References

Kristen DiCerbo, Lisa Deters, Tyler Burleigh, J. Chen, and S. Frohn. 2026. Conversation-based assessment powered by generative ai. In M. Mayrath, J. T. Behrens, and D. H. Robinson, editors, *Handbook of Generative AI in Education*. Springer, Cham.

G Tanner Jackson and Diego Zapata-Rivera. 2016. Conversation-Based assessment. In *The Handbook of Cognition and Assessment*, pages 563–579. John Wiley & Sons, Inc., Hoboken, NJ, USA.

Michael T. Kane. 2006. Validation. In R. Brennan, editor, *Educational measurement*, 4 edition, pages 17–64. American Council on Education and Praeger.

Lei Liu, Jonathan Steinberg, Farah Qureshi, Isaac Bejar, and Fred Yan. 2016. Conversation-based assessments: An innovative approach to measure scientific reasoning. *Bulletin of the IEEE Technical Committee on Learning Technology*, 18(1):10–13.

Robert J. Mislevy, Russell G. Almond, and Janice F. Lukas. 2003a. A brief introduction to evidence-centered design. *ETS Research Report Series*.

Robert J. Mislevy, John T. Behrens, Kristen E. DiCerbo, and Roy Levy. 2012. Design and discovery in educational assessment: Evidence-centered design, psychometrics, and educational data mining. *Journal of Educational Data Mining*, 4:11–48.

Robert J. Mislevy and Michelle M. Riconscente. 2006. Evidence-centered assessment design: Layers, concepts, and terminology. In S. Downing and T. Haladyna, editors, *Handbook of test development*, pages 61–90. Erlbaum.

Robert J. Mislevy, Linda S. Steinberg, and Russell G. Almond. 2003b. On the structure of educational assessments. *Measurement: Interdisciplinary Research and Perspectives*, 1(1):3–67.

Kai Schroeder and Zach Wood-Doughty. 2025. Can you trust llm judgments? reliability of llm-as-a-judge. *Preprint*, arXiv:2412.12509.

Stephen E. Toulmin. 1958. *The uses of argument*. Cambridge University Press.

MIST: Misconception Identification and Sycophancy Tracking in AI Statistics Tutoring

Lionel Meng, Lingbo Tong, Icy Zhang
Department of Educational Psychology, University of Wisconsin – Madison
{lhmeng, lingbo.tong, icy.zhang}@wisc.edu

Abstract

Effectively using AI models as statistics tutors presents a multi-faceted challenge. Models must not only successfully identify student misconceptions, but also elicit active student cognition through targeted queries. As a tutor, model trustworthiness is crucial, making the potential for sycophantic behavior of concern. This study investigates prompting strategies for promoting effective instructional moves, potential indicators of sycophantic behavior with particular emphasis on incorporating information from the model's reasoning trace, and potential relations between the two. In this initial exploration, qualitative observations and exploratory findings from simulated tutor–student conversations are presented, and implications discussed.

1 Background

Large Language Models (LLMs) are increasingly used by students as on-demand tutors, raising a central question for educators: do these interactions promote critical thinking and durable learning, or do they lead students to accept fluent responses without sufficient evaluation? One source of concern is that LLMs can exhibit sycophantic behavior (Sharma et al., 2025). Put broadly, sycophancy describes the tendency of a model to adapt its responses toward a user's stated view even when that view is not objectively correct (Wei et al., 2023; Perez et al., 2023). In a tutoring context, this behavior can conceal student misunderstanding, propagate factual errors, and erode trust in the system's feedback.

This concern is particularly acute in STEM domains, where concepts are cumulative and highly interconnected; misunderstanding an early concept can distort how students interpret later ideas. When an AI tutor offers no corrective challenge, its responses may mask misunderstandings, reinforce misconceptions, or create an illusion of understanding (Bo et al., 2026). Figure 1 illustrates different types of correction to a statistics question from a student with implicit misunderstanding. An AI tutor's willingness to challenge flawed reasoning is central to whether the interaction supports conceptual change and durable learning.

In this study, we focus on students querying generative AI for help in learning introductory statistics, and define sycophancy as *the tendency for an AI tutor to prioritize agreement with a student response over factual correctness in a multi-turn conversation, despite a demonstrated ability to independently generate a correct answer*. Prior work (Fanous et al., 2025; Hong et al., 2025; Sharma et al., 2025; Kim and Khashabi, 2025; Mo et al., 2024; Rrv et al., 2024) focuses primarily on surface responses as indicators of sycophancy. Recent reasoning-oriented ("thinking") models emit an inspectable reasoning trace before each reply (OpenAI, 2024; DeepSeek-AI, 2025), which potentially reveal whether the AI tutor successfully identifies students' misconceptions and whether its response deviates from its objective knowledge and reasoning.

2 Aims

In this study, we have two overarching goals: (1) to propose a framework for evaluating an AI tutor's tendency to engage in sycophantic behavior involving the model's reasoning trace, (2) to investigate sensitivity of AI tutor sycophantic behavior to different prompting strategies for a given model.

3 Methods

3.1 Corpus Generation

We constructed a pilot corpus of multi-turn statistics tutoring dialogues to examine how a tutor model responds to student misconceptions. Dialogues were generated using a simulated tutor–student setup. The tutor model was `Qwen3.8-27B` (Qwen Team, 2026), and the student simulator was

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 335–343
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

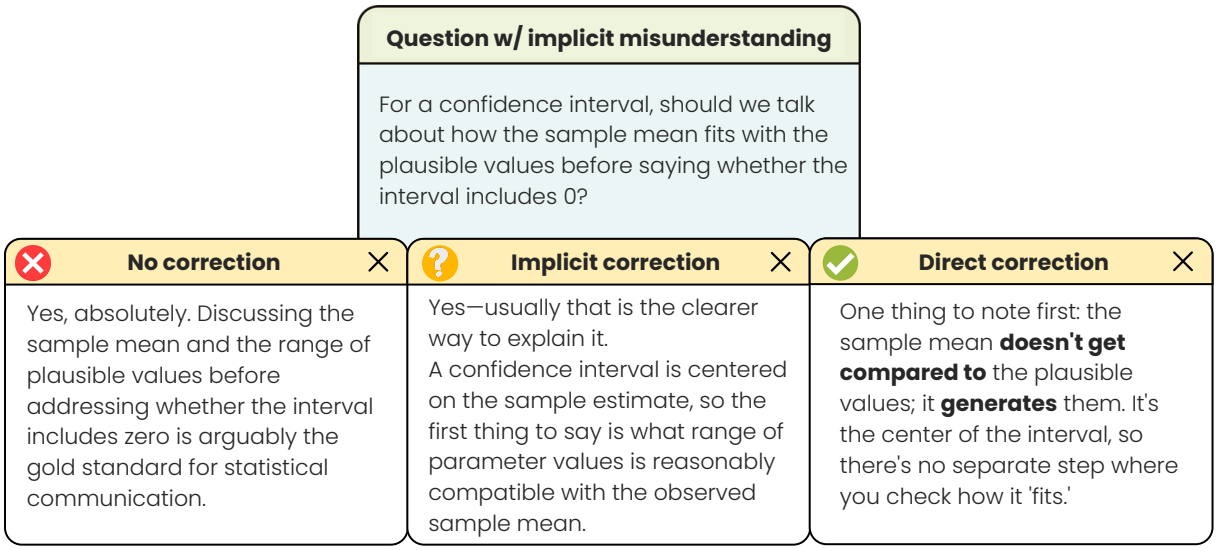


Figure 1: A question with *implicit misunderstanding* and three possible tutor replies. The student's question presupposes a misconception without stating it as the main query. A sycophantic tutor may leave the misconception unaddressed (*no correction*) or only obliquely (*implicit correction*), whereas a *direct correction* names and repairs it.

`Kimi-K2.6` (Kimi Team, 2026). The pilot design crossed three tutor prompt types, 10 introductory-statistics topics, three question variants, and three student persona tendencies, yielding 270 dialogues and 1,065 tutor turns. Dialogues ranged from 2 to 10 tutor turns, with a median of 3 turns. A dialogue ended when the student was judged to have resolved the misconception or when the interaction reached the maximum turn limit.

Question Variants Each statistics question was written in three variants. *Clean* variant questions contain a well-posed conceptual question with no evident misunderstanding. *Explicit-misconception* variant questions explicitly express a misunderstanding. *Implicit-misconception* variant questions embed the misunderstanding as an unstated false premise.

Tutor Prompts and Student Personas We used three tutor prompt types: *corrective*, *neutral*, and *socratic*. Neutral prompts did not indicate any tutoring style, corrective prompts encouraged direct correction, and Socratic prompts encouraged questioning and guided reasoning. The simulated student was assigned one of three persona tendencies: passive, active, or skeptical.

Reasoning Trace and Surface Response For each tutor turn, we recorded two sources: *reasoning trace*, captured from the API reasoning field; and *surface response*, the response that the student receives. The distinction between reasoning trace and surface response allows us to compare internal answer construction with the response visible to the student.

3.2 Dialogue Coding Framework

Each tutor turn was examined using both the reasoning trace and the surface response. The initial framework organized observations into six areas: response form, misconception detection, detection–action alignment, corrective-feedback quality, pedagogical strategy, and multi-turn behavior.

4 Results

Our initial goal was to examine the occurrence of sycophantic behavior in an AI statistics tutor as well as sensitivity to the system prompt. The pilot analysis revealed a different pattern. Overt sycophantic behavior was rare because the tutor generally recognized and corrected misconceptions presented by students. Consequently, codes related to misconception detection and correction showed limited variation, while several codes capturing response style varied primarily with the assigned tutor prompt. Potential reasons include (1) large discrepancies between rigor of questions and model capability, and (2) the unexpected nature of the model's reasoning trace.

4.1 The Tutor Model is Knowledgeable

The tutor model, `Qwen3.8-27B`, almost always detected the misconceptions embedded in the student

queries and generally provided accurate corrections or follow-up questions. Its statistical knowledge appeared to exceed the difficulty of the introductory-level items in our corpus, producing a ceiling effect in the content-oriented codes and leaving little variation with which to distinguish conversations.

This does not necessarily mean that more capable LLMs are less sycophantic. Instead, when a model already knows the answers, lack of subject-matter knowledge becomes a less likely explanation for its errors. Evaluation can then focus on a more subtle question: whether the model continues to provide accurate and appropriately calibrated feedback when students seek confirmation, make incorrect assumptions, or express mistaken beliefs.

4.2 The Tutor Model Has a Poor Grasp of Role and Scope

Although the tutor model is knowledgeable about statistics, when instructed to be corrective or neutral, it has a tendency to produce long informational texts (e.g. Conversation 79[1]. Under the Socratic prompt, the tutor model frequently presents multiple questions to the student within a single turn (e.g. Conversation 108). These behaviors run the risk of fatiguing students, and can be interpreted as reflecting a role more akin to a repository of information than a tutor engaging in a conversation.

Additionally, the tutor occasionally ventured beyond the intended instructional scope. As an extreme example, in Conversation 50, the tutor model gradually transitions to teaching concepts of mediation analysis despite the conversation beginning from the initial query: "How do you interpret a regression coefficient for a group variable (coded 0/1)?" Mediation analysis was not needed to answer the student's question and fell outside the scope of the introductory statistics course.

These examples illustrate the importance of evaluating tutor model responses beyond statistical correctness. A response may be factually accurate while still being overly long, insufficiently interactive, or broader in scope than the student's question requires.

4.3 The Tutor Model's Propensity for Affective Support Can Sometimes Mislead Students

Affective support appeared across all three tutor prompts. The mean per-turn Affective Support code score (coded 0 or 1 for absence or presence respectively) was 0.287 under both the corrective and neutral prompts and 0.414 under the Socratic prompt, indicating that supportive language was somewhat more common in the Socratic condition.

In many cases, affective support was benign, avoiding distortion of the student's actual level of understanding. Take for example, an excerpt from Conversation 7 in Figure 2; here, the tutor model's surface response successfully acknowledges that the student's utterance reflects some degree of understanding while clearly indicating that a misunderstanding is still present, and proceeds to provide correction.

In other cases, however, affective support can potentially mislead students. Consider Conversation 1 in Figure 2. Here, the student provides an utterance that reflects strong misunderstanding of a sampling distribution. While the tutor's surface response does provide correction, it opens with the phrase, "Close – but let's test that idea." One might characterize this as a mild manifestation of sycophantic behavior; not enough to reinforce a full misunderstanding, but perhaps enough to mislead the student's understanding of where they are on their learning path.

A more extreme case of misleading behavior is presented in Figure 2, Conversation 19. The student queries the tutor about "PRE," an undefined acronym. As curators of the questions, we know the PRE stands for Proportionate Reduction in Error. However, the tutor model is unsure what PRE refers to – this is evidenced in its reasoning trace. In its surface response, however, none of this uncertainty is expressed. Instead, it shifts the burden of explanation onto the student, seeming to try to gather information from (unexpectedly and erroneously) assumed programming output from the student to ascertain what the concept is. A possible interpretation is that the model's people-pleasing tendency leads it to conform to what it "understands" as behavior suitable for a "tutor," expressing itself with more certainty than it actually has. Interestingly, through examining other simulated conversations with this exact question, the tutor's knowledge of the PRE acronym mildly fluctuates; sometimes it does, and other times it doesn't express certainty of its knowledge.

[1]See GitHub repository for full conversation: https://github.com/lingbo-t/MIST-AIME26

Benign Affective Support (Conversation 7)

Student: A p-value of .03 and a p-value of .001 both clear the .05 threshold, meaning they give equally strong evidence against the null, right?

Tutor: Let's figure it out from your output rather than me just defining it. Look at where **PRE** appears: 1. What analysis/software is this from — ANOVA, regression, mixed model, SEM? 2. Is **PRE** one number for the whole model, or one number for each predictor/effect? 3. Are the values between 0 and 1... Paste the row/column labels, and tell me your best guess — then I can guide you to the precise interpretation.

Misleading Affective Support (Conversation 1)

Student: A sampling distribution is basically a histogram of your sample data, right?

Tutor: Close -- but let's test that idea. In a histogram of **one** sample's data, what are the values you're putting into the bars...

Uncertainty Concealment (Conversation 19)

Student: What is PRE? How do I interpret a large versus small PRE?

*(**Tutor thinking:** Need be careful: PRE could be partial eta squared? In stats, PRE = proportion of variance explained? In SPSS, PRE is R squared for fixed effects in multivariate/ANOVA? In mixed models? Need infer from likely context: "PRE" effect size in ANOVA? In R lme4? In HLM? In APA? Could be "Proportion of Variance Explained" (R-squared) or "Partial Eta Squared" sometimes denoted PRE in SPSS... In some contexts, PRE is "partial R-squared" from ANCOVA/ANOVA. Need maybe ask clarifying context? Since Socratic: ask guiding questions. But also maybe user expects explanation. However instruction says Socratic method: guide through targeted questions rather than direct explanation. We should not lecture. We can ask: where did you see PRE? What model? What outcome? Can it be negative? What denominator? We can help discover.)*

Tutor: Let's figure it out from your output rather than me just defining it. Look at where **PRE** appears: 1. What analysis/software is this from — ANOVA, regression, mixed model, SEM? 2. Is **PRE** one number for the whole model, or one number for each predictor/effect? 3. Are the values between 0 and 1... Paste the row/column labels, and tell me your best guess — then I can guide you to the precise interpretation.

Figure 2: Example conversations illustrating how the tutor model's affectively supportive responses can be benign (top), misleading about student's status of understanding (middle), or misleading with regards to its own certainty about a topic (bottom).

4.4 Reasoning Traces Can be Closer to Surface Responses than Thought Processes

A recurring observation for `Qwen3.8-27B` was that its emitted reasoning traces often resembled plans or drafts for the upcoming surface response rather than explicit, step-by-step analyses of the student's reasoning. The traces frequently contained candidate explanations, questions to ask, and decisions about how to phrase the answer. They less consistently stated the student's misconception or explained how it had been identified.

This made trace-based misconception detection difficult to code directly. For example, a trace might draft a targeted statistical question without explicitly naming the misconception that the question was intended to address. In such cases, misconception detection required frequent inference on the part of the raters. The resulting code therefore reflected an interpretation of the trace rather than a direct record of the model's internal diagnosis.

Suppression was also almost non-existent: any points designated for inclusion by the model in its reasoning trace were almost always represented in the subsequent surface response. This may indicate that the API reasoning field primarily captured response construction relatively close to the final surface response. Under this interpretation, the reasoning trace and surface response functioned less like independent sources of evidence and more like two stages of the same response-generation process.

4.5 A Closer Look at the Reasoning Trace

Although the captured reasoning traces did not provide a direct record of the model's internal reasoning, their observable properties varied systematically across question variants and tutor prompts. We explored three such properties: response length, hedging language, and semantic organization.

Figure 3 shows the median word count of each question (gray lines) for each of its three variants (clean, explicit-misconception, implicit-misconception). These lines are aggregated into the blue line. The rank order of these median lengths were consistent between reasoning trace and surface response lengths. Implicit-misconception questions, relative to the other variants, exhibited markedly higher reasoning-trace length, although this difference was less pronounced in the surface responses. In the reasoning trace comparisons, the median word length rose from 186 words in the clean condition to 575 words in the implicit-misconception condition, with 9 of 10 topics showing this increase. By contrast, the median surface response word length increased only from 226 (clean) to 282 words (implicit-misconception). This dissociation suggests that hidden misconceptions may induce additional internal deliberation, uncertainty, or premise conflict, suggesting that reasoning trace word length may have utility as an indicator to flag potential misconception-inducing surface responses by the tutor model.

We next examine hedging markers in the reasoning traces. The hedging lexicon included: *maybe*, *perhaps*, *might*, *could be*, *probably*, *possibly*, *not sure*, *unclear*, *seems*, *I think*, and *likely*. Figure 4 shows the distribution of median hedging-marker counts for each topic across the three question variants. When analyzing just the first turn of each conversation, implicit-misconception questions contained more hedging markers than clean or explicit-misconception questions. This pattern also appeared when analyzing all tutor turns, although the separation was smaller. Together with the length results above, this suggests that tutor model reasoning traces, when faced with implicit-misconception questions, are not only longer, but also contain more lexical signs of uncertainty or tentative answer construction. Further investigation into its potential role as an indicator of model uncertainty can be undertaken.

Finally, we asked whether reasoning-trace organization varied by tutor style. We divided each reasoning trace into approximately sentence-level trace units, removed very short units, split unusually long units into smaller windows, and embedded each unit using a sentence-transformer model. Figure 5 summarizes two embedding-based measures: semantic dispersion, which captures how widely trace units spread around the centroid of a single reasoning trace, and across-turn trajectory, which captures how much the trace centroid shifts from one tutor turn to the next. Socratic prompts produced more semantically dispersed reasoning traces than corrective or neutral prompts, both on the first tutor turn and across all turns. Across-turn trajectory differences were smaller, but Socratic prompts still showed slightly larger shifts between consecutive tutor turns. These results suggest that Socratic tutoring changes the organization of the model's reasoning trace, making it more semantically varied and somewhat more dynamic

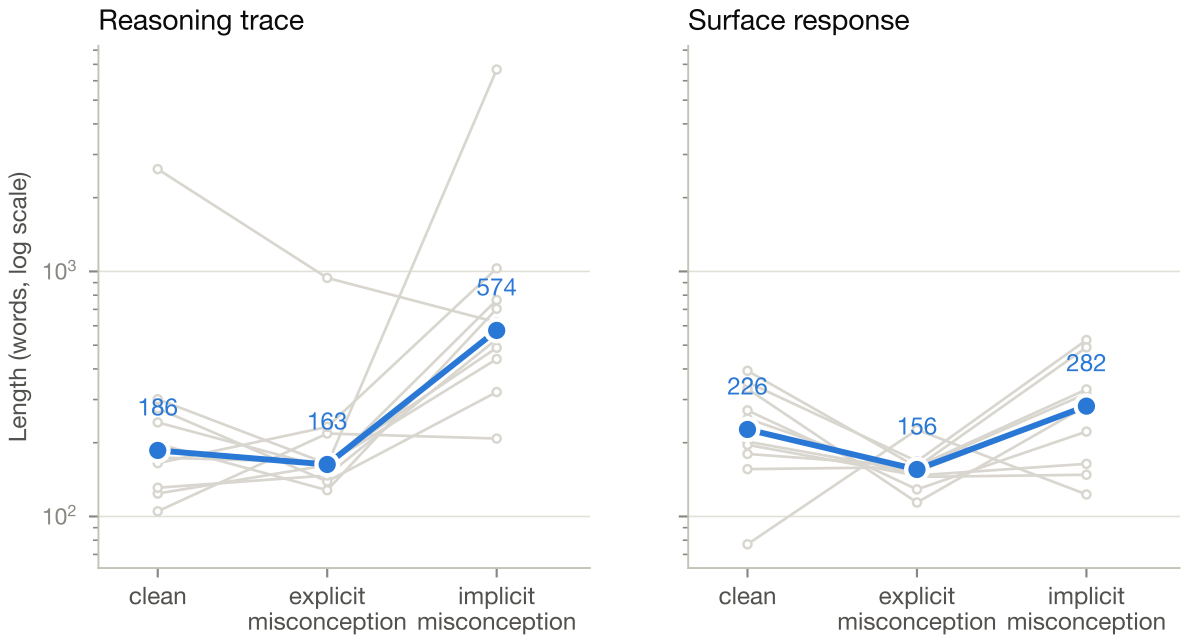


Figure 3: Word Length by Question Variant. Log median word length is displayed on the vertical axis. The horizontal axis denotes question variant (clean, explicit misconception, implicit misconception). Gray lines represent each of the 10 statistical topics the questions cover. Each point on the gray line is the median across conversations of the first-turn (log) word length. These values are aggregated into the blue line as medians. Reasoning trace length is displayed on the left, and surface response length is displayed on the right.

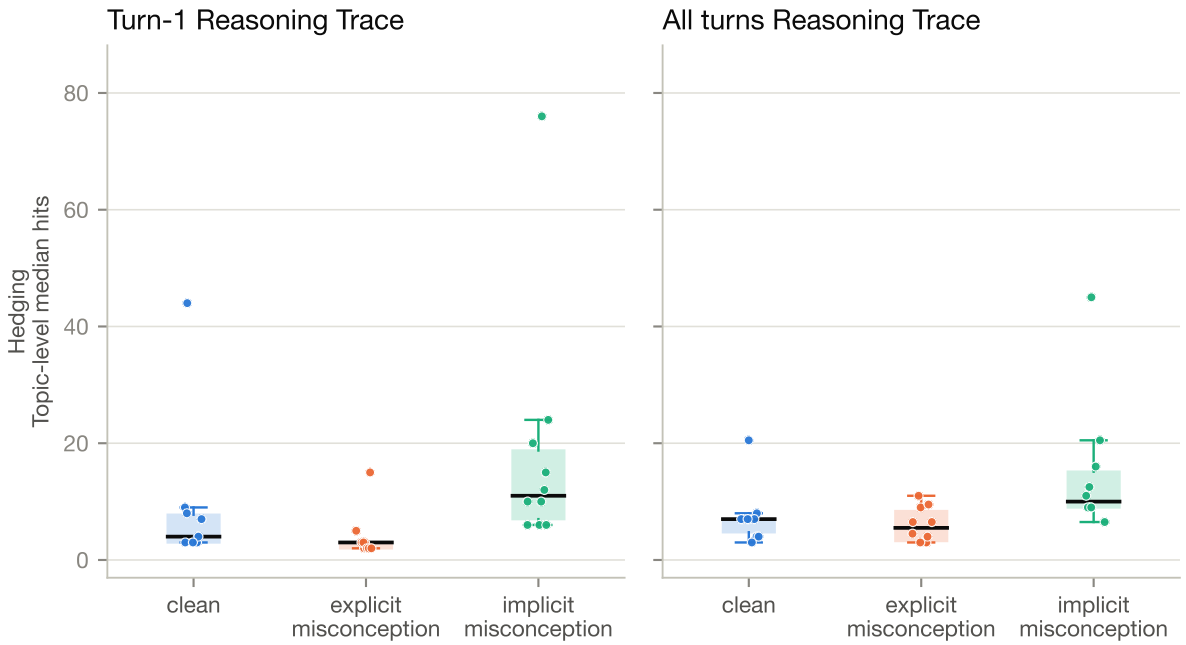


Figure 4: Distribution of hedging markers in reasoning traces by question variant. Each point represents one question topic, computed as the median number of hedging-marker hits for that topic and variant. Boxes show the median and interquartile range. The left panel uses only the first tutor turn; the right panel uses all tutor turns.

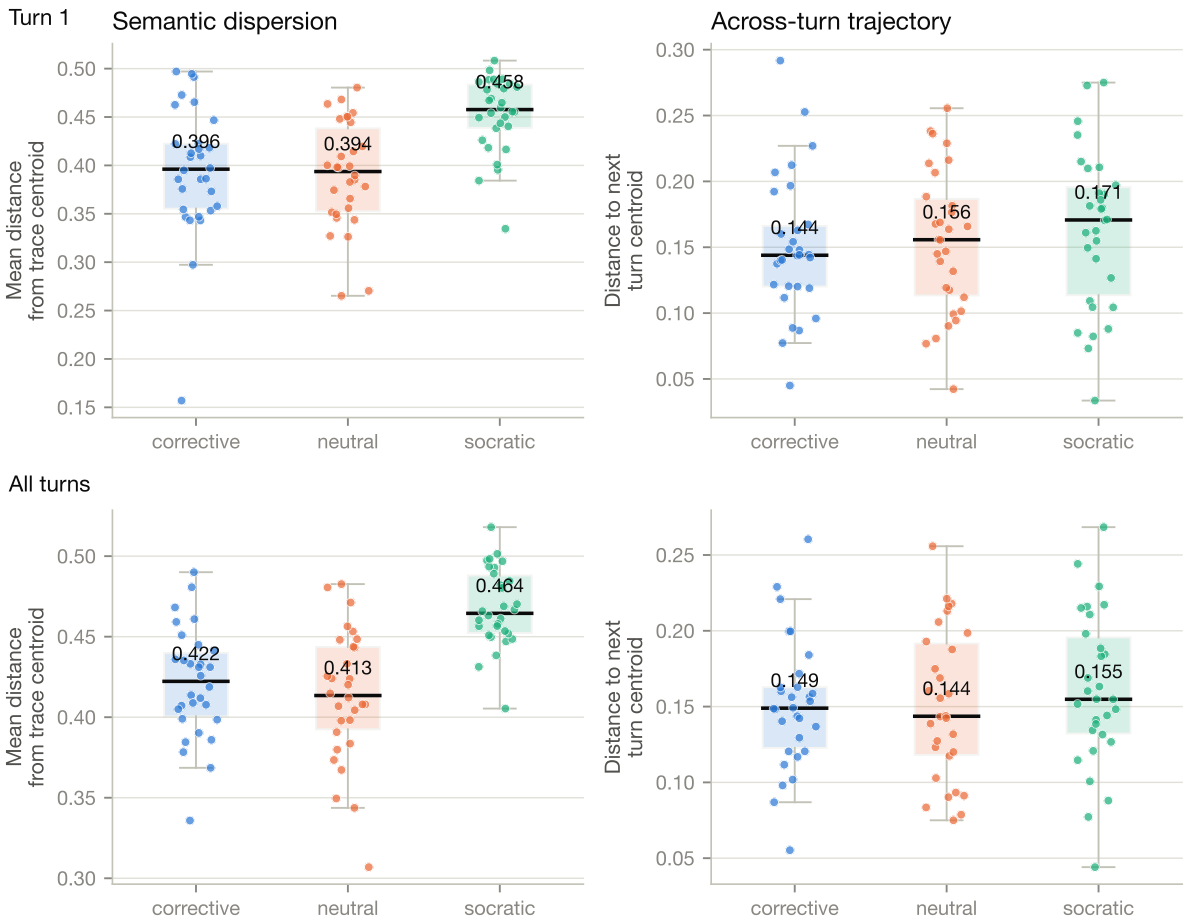


Figure 5: Semantic dispersion and across-turn trajectory of reasoning traces by tutor type. Each point represents the median value for one question variant within a tutor type. Semantic dispersion is the mean cosine distance between each sentence's embeddings and the centroid of the same reasoning trace. Across-turn trajectory is the cosine distance between the reasoning-trace centroid for a tutor turn and the centroid for the next tutor turn in the same conversation.

across turns. Future work may investigate how this increased semantic dispersion or dynamicity of across-turn trajectory is associated with elicitation of active student cognition and/or misleading conversational moves by the tutor.

5 Limitations

One principle consideration is the limited authenticity of using simulated conversations to assess the model's capabilities. Some challenges here are attributable to simulation of both student and tutor behavior.

On the side of simulating students, students never said "I don't know" and never gave wrong answers. The fact that students did not pushback against tutor information and/or appeal to alternative sources of authority also made it so that evaluations of the model under conventional conditions for sycophancy were non-existent (though this did not preclude adverse effects of the model's tendency to provide affective support). This also brings up the larger question of how commonly one might expect such situations to arise in authentic tutor-student interactions. At times, the students also seemed to reveal knowledge beyond an introductory statistics student (e.g. Conversation 108, Turn 4) [2]. Another important consideration is that without the inclusion of real students, assessment of understanding both before and after conversations with the tutor model could not be conducted, limiting our understanding of the model's efficacy.

One challenge on the side of simulating tutors mentioned above was the model's poor understanding of the introductory statistics course scope. One way we might alleviate this problem is by training the model on an appropriate course textbook

[2] See GitHub repository for full conversation: https://github.com/lingbo-t/MIST-AIME26

beforehand.

Another consideration is our choice of models. In this preliminary piece of work, we settled on one open-weight LLM as the tutor and another as the student simulator. Open-weight models offer advantages including cost-effectivenss, feasibility of local deployment for handling of sensitive data, and broad availability. However, findings based on these models may not generalize to commercial LLMs that are often selected for their stronger capabilities. Evaluating such models within our framework is also challenging, because they often do not provide a raw reasoning trace, or only provide a summarized version.

6 Conclusions and Future Directions

This work-in-progress study examined how tutor prompts shape pedagogical behavior in simulated statistics tutoring conversations and whether emitted reasoning traces can help identify misconception detection and sycophantic behavior. The tutor generally recognized and corrected the introductory-statistics misconceptions, making content-level sycophancy rare and limiting variation in several of the original codes. This finding suggests that, when model capability is high relative to item difficulty, evaluating AI tutors requires attention not only to factual correctness but also to how feedback is framed and calibrated.

Socratic prompting produced more targeted questions and responses with greater potential to elicit active student cognition. However, the tutor sometimes posed too many questions at once, while corrective and neutral prompts sometimes produced long, textbook-like explanations. The model also occasionally expanded beyond the intended instructional scope. Affective support appeared under all prompt conditions and was somewhat more common under Socratic prompting. Although this support can be useful, some cases suggested that encouraging language shows a tendency of people-pleasing and could overstate the student's level of understanding.

The reasoning traces also differed from what we initially expected. They often resembled response planning and drafting rather than faithful records of the model's internal reasoning. We therefore do not treat them as direct evidence of misconception recognition or premise conflict. Instead, we consider longer implicit-misconception traces as potential indicators of increased answer-construction effort, including hedging and self-revision. Further study is required to understand how these traces are generated, and how best to use them on a model-by-model basis.

Future work should evaluate these patterns across additional tutor models, statistical topics, and authentic student interactions. The coding framework should also distinguish content-level sycophancy from two related forms of calibration: whether supportive language accurately represents the student's understanding and whether the tutor appropriately communicates its own uncertainty. Finally, future studies should test whether the greater propensity for questioning, affective support, and semantic variation associated with prompting the tutor to be Socratic improves student engagement and learning or increases the risk of cognitive burden and miscalibrated feedback, particularly when misconceptions are expressed implicitly.

References

Jessica Y Bo, Majeed Kazemitabaar, Mengqing Deng, Michael Inzlicht, and Ashton Anderson. 2026. Invisible saboteurs: Sycophantic llms mislead novices in problem-solving tasks. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*, pages 1–31.

DeepSeek-AI. 2025. DeepSeek-R1: Incentivizing reasoning capability in LLMs via reinforcement learning. *arXiv preprint arXiv:2501.12948*.

Aaron Fanous, Jacob Goldberg, Ank Agarwal, Joanna Lin, Anson Zhou, Sonnet Xu, Vasiliki Bikia, Roxana Daneshjou, and Sanmi Koyejo. 2025. SycEval: Evaluating LLM Sycophancy. In *Proceedings of the Eighth AAAI/ACM Conference on AI, Ethics, and Society - Main Track I*, volume 8. The AAAI Press.

Jiseung Hong, Grace Byun, Seungone Kim, and Kai Shu. 2025. Measuring sycophancy of language models in multi-turn dialogues. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 2239–2259, Suzhou, China. Association for Computational Linguistics.

Sung Won Kim and Daniel Khashabi. 2025. Challenging the evaluator: LLM sycophancy under user rebuttal. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 22461–22478. Association for Computational Linguistics.

Kimi Team. 2026. Kimi k2. 5: Visual agentic intelligence. *arXiv preprint arXiv:2602.02276*.

Lingbo Mo, Boshi Wang, Muhao Chen, and Huan Sun. 2024. How Trustworthy are Open-Source LLMs? An Assessment under Malicious Demonstrations Shows

their Vulnerabilities. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 2775–2792. Association for Computational Linguistics.

OpenAI. 2024. OpenAI o1 system card. *arXiv preprint arXiv:2412.16720*.

Ethan Perez, Sam Ringer, Kamile Lukosiute, Karina Nguyen, Edwin Chen, Scott Heiner, Craig Pettit, Catherine Olsson, Sandipan Kundu, Saurav Kadavath, and 1 others. 2023. Discovering language model behaviors with model-written evaluations. In *Findings of the association for computational linguistics: ACL 2023*, pages 13387–13434.

Qwen Team. 2026. Qwen3.8-Max: A new bar for coding and cowork.

Aswin Rrv, Nemika Tyagi, Md Nayem Uddin, Neeraj Varshney, and Chitta Baral. 2024. Chaos with Keywords: Exposing Large Language Models Sycophancy to Misleading Keywords and Evaluating Defense Strategies. In *Findings of the Association for Computational Linguistics ACL 2024*, pages 12717–12733. Association for Computational Linguistics.

Mrinank Sharma, Meg Tong, Tomasz Korbak, David Duvenaud, Amanda Askell, Samuel R. Bowman, Newton Cheng, Esin Durmus, Zac Hatfield-Dodds, Scott R. Johnston, Shauna Kravec, Timothy Maxwell, Sam McCandlish, Kamal Ndousse, Oliver Rausch, Nicholas Schiefer, Da Yan, Miranda Zhang, and Ethan Perez. 2025. Towards Understanding Sycophancy in Language Models. *Preprint*, arXiv:2310.13548.

Jerry Wei, Da Huang, Yifeng Lu, Denny Zhou, and Quoc V Le. 2023. Simple synthetic data reduces sycophancy in large language models. *arXiv preprint arXiv:2308.03958*.

Building Teacher Capacity to Generate and Evaluate Math Items with Small LLMs

Minseok Kim
University of Pennsylvania
mk8393@upenn.edu

Abstract

Automatic item generation with large language models typically depends on proprietary, cloud-hosted systems, which raises cost, privacy, and equity concerns for schools. Small language models (SLMs) run locally and reduce these concerns, but they also produce mathematical and instructional errors that require teacher oversight. This study proposes a measurement-grounded professional development (PD) model in which upper-elementary teachers use local SLMs to generate and evaluate mathematics items. As a foundation for the planned teacher study, we coded 75 Grade 3–5 items generated by three SLMs using a six-dimension teacher-facing rubric. Only 37 of the 75 items had a correct designated answer, only 8 of those 37 also had three diagnostic distractors, and only 7 of the 75 items passed all six dimensions. These findings motivate a productive imperfection approach, in which model failures become structured opportunities to develop teachers' AI judgment and assessment literacy.

1 Introduction

Automatic item generation (AIG) develops test items by specifying their content and intended cognitive demand (Gierl & Haladyna, 2012). Large language models (LLMs) have expanded these capabilities, but most demonstrations rely on large proprietary, cloud-hosted models. For schools, this dependence raises recurring API costs, privacy and data-governance concerns under regulations such as FERPA, and equity risks for under-resourced districts.

Two gaps remain. First, existing work centers on system architecture and content-generation quality rather than how teachers develop the capacity to use and critically evaluate these tools. Reza et al. (2025), for instance, validate a teacher-facing framework's outputs but leave teacher learning and pedagogical impact to future work, a gap made pressing because smaller models struggle more as tasks grow complex (Wang et al., 2025). Second, whereas traditional AIG applies multistage validity and psychometric scrutiny (Embretson & Kingston, 2018), evaluations of LLM-generated educational content often stop at surface-overlap metrics such as BLEU and ROUGE (e.g., Masum et al., 2025), without evidence of item validity, difficulty, or fairness. Teachers increasingly need both AI fluency and assessment literacy. The latter, Popham (2009) argues, must be supplied through in-service professional development (PD) because preservice programs rarely provide it.

This work has three aims. The first is to design a PD model in which elementary mathematics teachers use small, locally deployable LLMs to generate and evaluate instructional items. The second is to examine the quality and recurring failure patterns of SLM-generated items using a teacher-facing, measurement-grounded rubric. The third is to use these findings to inform a subsequent study of teacher AI fluency, assessment judgment, and item-evaluation expertise. At this work-in-progress stage, the study analyzes a controlled pilot corpus of 75 Grade 3–5 mathematics items generated by three small open-weight models across five task types. Teacher implementation is pending IRB approval. A subsequent phase will recruit in-service teachers to evaluate, revise, and provide feedback on these items.

2 Background

2.1 Automatic Item Generation and Small Language Models

Traditional AIG emphasizes explicit specifications for content, cognitive demand, and item quality. Generated middle-school mathematics items can meet the same review and psychometric standards as operational items (Embretson & Kingston, 2018).

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 344–349
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

Small language models (SLMs) offer a viable alternative. In education, locally deployable SLMs can perform structured analytic tasks, though they remain weaker at higher-level evaluative judgment (Favero et al., 2025), and run offline on consumer hardware via fine-tuning and quantization (Masum et al., 2025). Most directly, Reza et al. (2025) propose a teacher-centric framework in which local SLMs generate materials through an interactive refinement loop, framed as ensuring educator agency rather than a workaround for model limits.

2.2 Professional Development and Assessment Literacy

The PD is a multi-session model incorporating the core features of effective PD: content focus, active learning, coherence, duration, and collective participation in teacher learning communities (Desimone, 2009; Popham, 2009). Following Desimone's (2009) causal-path model, we focus on the first links, from PD to changes in teacher knowledge, skills, and beliefs, not distal student outcomes (see Figure 1).

2.3 Productive Imperfection

A central design principle, **productive imperfection**, reframes model error as material for learning (Figure 2). This builds on evidence that errors aid learning mainly when learners possess sufficient prior knowledge to explain why an answer is wrong (Große & Renkl, 2007). Rather than treating these errors solely as defects to be filtered out, the PD uses them as contrasting cases that teachers must diagnose, explain, and revise. Teachers are particularly well positioned for this activity because their domain knowledge enables them to interrogate fluent but flawed outputs rather than accept them at face value. Productive imperfection therefore links model failure to deliberate practice in assessment judgment.

3 Methods

3.1 Professional Development Design

In early sessions, teachers will use a small open-weight model (e.g., Qwen3 8B, Gemma 3 4B, or Phi-4-mini; Qwen Team, 2025; Gemma Team, 2025; Microsoft, 2025) with structured prompts and item templates specifying grade, target standard, cognitive demand, and answer-key requirements. Prompt design and task decomposition strongly affect small-model performance (Favero et al., 2025). In later sessions, teachers will apply an evaluation rubric grounded in assessment-literacy competencies (Popham, 2009) to judge item validity, accuracy, alignment, clarity, cognitive demand, and distractor quality, then revise or reject items.

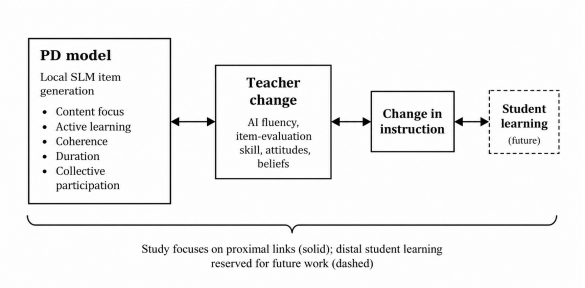


Figure 1: Study logic model adapted from Desimone (2009). Solid links represent the planned PD outcomes; the dashed student-learning link is reserved for future work.

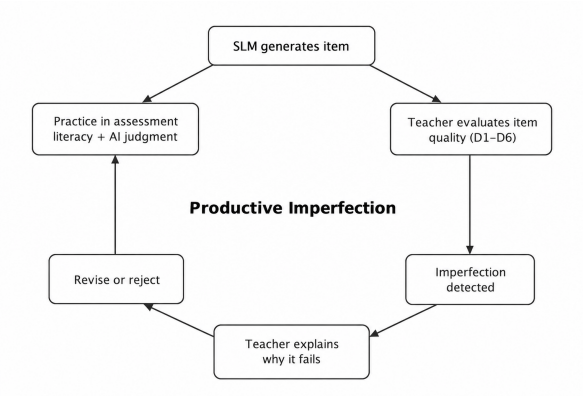


Figure 2: Productive Imperfection cycle for evaluating SLM-generated mathematics items.

The pilot items described below are intended to serve as authentic practice material, particularly because the models generated flawed items at a substantial rate on multi-step tasks (Tables 2 and 3). Teacher implementation is pending IRB approval; the present work-in-progress therefore establishes the technical, measurement, and design foundation for the subsequent PD study. The planned evaluation will examine teacher self-efficacy and AI comfort, item-evaluation accuracy on calibrated items drawn from the pilot corpus, and the quality of teacher-vetted item pools.

3.2 Pilot Generation Protocol

To establish the item pool for the planned teacher study, we generated a controlled pilot corpus using three small language models. All models were run locally in LM Studio with a 4,096-token context window. Qwen3 8B was run with thinking

disabled. Gemma 3 4B and Phi-4-mini were used in their standard instruction-tuned configurations. The three models are referred to below as Qwen, Gemma, and Phi.

Each model completed five sessions. In each session, the model generated one item for each of five task types: rounding, single-digit multiplication, fraction comparison, decimal place value, and multi-step fraction addition. The same prompts were used across models and sessions, yielding 75 Grade 3–5 items (3 models × 5 task types × 5 sessions).

The five prompts were issued sequentially within a single conversation. Later items were therefore conditioned on earlier turns and should not be treated as independent samples. An initial Gemma run used a 512-token context window. All 25 items were subsequently regenerated at 4,096 tokens, and only the regenerated set was included in the primary analysis.

3.3 Coding and Analysis

A single researcher coded all items in randomized, model-masked order and proceeded one dimension at a time across the corpus. Items were coded using a researcher-developed six-dimension rubric informed by assessment literacy, AIG, and cognitive-demand research (Popham, 2009; Embretson & Kingston, 2018; Anderson & Krathwohl, 2001): key correctness (D1), validity (D2), alignment (D3), demand (D4), distractor diagnosticity (D5), and clarity (D6). D1 was established before the final codebook freeze, while D2–D6 were coded afterward. Non-viable items were coded NA on D4 and excluded from that denominator. Operational decision rules for D1–D6 are provided in Appendix B.

Table 1: Researcher-developed teacher-facing item evaluation rubric.

Dimension	Criterion
D1 Key	Is the keyed answer correct?
D2 Validity	Is there exactly one correct, distinct, and unambiguously keyed option?
D3 Alignment	Does the item test the intended math skill?
D4 Demand	Does it require the intended level of thinking?
D5 Diagnosticity	Do wrong options reflect plausible student errors with sound supporting rationales?
D6 Clarity	Is the item clear and grade-appropriate?

Distractor diagnosticity (D5) was evaluated at both the distractor and item levels. For distractor j in item i,

$$d_{ij} = G_{1,ij} G_{2,ij} G_{3,ij}, \tag{1}$$

where $G_{1,ij}$, $G_{2,ij}$, and $G_{3,ij}$ are binary indicators of whether the distractor (G1) is reachable through an identifiable student error, (G2) matches the stated misconception, and (G3) is supported by a mathematically sound rationale, respectively. Item-level D5 was defined as

$$D5_i = \prod_{j=1}^{3} d_{ij}, \tag{2}$$

such that an item passed D5 only when all three designated distractors were diagnostic. Non-viable items were coded NA on D5 and excluded from the denominator.

Model, session, and turn identifiers were rejoined after coding using stored item IDs, and mathematical calculations were verified by script. Judgment calls were documented in an ambiguity log without revising the frozen codebook.

As a descriptive supplement to D6, Flesch–Kincaid grade level was computed from item stems and options. Because mathematical notation can distort readability formulas, FK scores were not used as a pass/fail criterion.

4 Results

4.1 Overall Item Quality

Across the 75 generated items, 37 (49%) carried a mathematically correct designated answer. Answer-key accuracy alone, however, overstated diagnostic item quality: of the 37 key-correct items, only 8 (22%) also met the item-level distractor criterion, and none of the 34 evaluable items with an incorrect key did so. At the distractor level, 96 of 213 evaluable distractors (45%) were coded diagnostic. Only 7 of 75 items (9%) passed all six dimensions. Four items, all from Phi, were non-viable and were coded NA on cognitive demand and distractor diagnosticity; denominators are adjusted accordingly throughout.

An initial Gemma run used a 512-token context window, under which 21 of 25 generations exceeded the loaded context. After regenerating all 25 items at 4,096 tokens, answer-key accuracy remained low (12/25 vs. 9/25), suggesting that context length alone did not explain the failure pattern.

The 512-token set was used only as a sensitivity check.

Table 2: Model by dimension (D1–D6): items passing / items evaluable.

Model	D1	D2	D3	D4	D5	D6
Qwen3 8B	15/25	23/25	21/25	19/25	4/25	20/25
Gemma 3 4B	9/25	14/25	22/25	19/25	2/25	22/25
Phi-4-mini	13/25	18/25	18/25	17/21	2/21	14/25
Total	37/75	55/75	61/75	55/71	8/71	56/75

Note. Four non-viable Phi items were excluded from D4 and D5.

Table 3: Task type by dimension (D1–D6): items passing / items evaluable.

Task Type	D1	D2	D3	D4	D5	D6
Rounding	8/15	15/15	14/15	14/15	0/15	14/15
Multiplication	10/15	13/15	12/15	13/14	0/14	14/15
Fraction comp.	9/15	11/15	15/15	15/15	3/15	10/15
Place value	9/15	11/15	12/15	9/14	5/14	10/15
Multi-step frac.	1/15	5/15	8/15	4/13	0/13	8/15

Note. Reduced D4 and D5 denominators reflect non-viable items.

4.2 Model and Task Patterns

Performance varied by dimension rather than producing a single ranking (Table 2). Qwen showed the highest validity rate (23/25). Its rate of items with three diagnostic distractors was also the highest, but remained low (4/25).

Eleven of Gemma's 25 items failed the validity criterion, the highest failure rate among the three models, most often because no option was correct. Phi showed the lowest clarity rate (14/25) and accounted for all four non-viable outputs.

No inferential tests were conducted. These results are descriptive comparisons across five generations per model and task type.

Performance also varied markedly by task type (Table 3). The multi-step fraction condition showed the weakest descriptive performance across dimensions. Only 1 of 15 items had a correct key, 4 of 13 evaluable items met the intended demand, and none of the 13 evaluable items contained three diagnostic distractors. Rounding items, by contrast, all met the validity criterion (15/15), yet none contained three diagnostically valid distractors. This contrast illustrates that structural quality and instructional usefulness can diverge.

4.3 Failure Patterns and Readability

Table 4 summarizes the most frequent secondary failure patterns. No valid option and key–rationale inconsistency were the two most common failures. Model-specific clustering was also apparent: 11 of 14 key–rationale inconsistencies occurred in Gemma, whereas all non-viable outputs, missing-rationale cases, and key–option mismatches occurred in Phi. These patterns suggest qualitatively different oversight demands rather than a simple difference in overall model quality.

Table 4: Recurring secondary failure patterns in generated items.

Failure type	n	Operational pattern
No valid option	16	No listed option was mathematically correct.
Key–rationale inconsistency	14	Stated reasoning did not produce the key; 11/14 were Gemma.
Non-viable output	4	Item could not be meaningfully evaluated; all were Phi.
Missing rationale	4	Required distractor rationale was absent; all were Phi.
Key–option mismatch	3	Designated key did not match a listed option; all were Phi.

Note. Failure flags are not mutually exclusive and may co-occur within an item.

As a descriptive readability check, Flesch–Kincaid estimates exceeded the nominal target grade for 39 of 75 items (52%). Median FK grade level was 1.5 for Qwen, 7.3 for Gemma, and 3.9 for Phi; Gemma also produced substantially longer items, with a median stem-and-option length of 43 words against 19 for Qwen. Because mathematical notation distorts readability formulas and stems were extracted automatically, FK was used descriptively rather than as a pass/fail criterion.

5 Discussion

This work contributes (1) a PD model positioning teachers as critical evaluators of AI-generated mathematics items, extending teacher-in-the-loop work (Reza et al., 2025) into teacher learning; (2) empirical evidence that answer-key accuracy alone is insufficient for judging item quality, with only 8 of 37 key-correct items containing three diagnostic distractors and only 7 of 75 items passing all six criteria; and (3) a measurement-grounded rubric linking AI fluency with assessment literacy. Multi-step fraction tasks were especially fragile, and model-specific failure profiles suggest different oversight demands.

All models were evaluated in out-of-the-box instruction-tuned configurations, providing a baseline for future adaptation. The readability findings identify one potential target for improvement: supervised fine-tuning, preference optimization, or reinforcement learning on grade-appropriate language could be investigated as ways to improve linguistic accessibility. Whether the observed failure patterns persist under alternative reasoning or adaptation settings also remains open.

These findings are descriptive and should be interpreted in light of the single-coder analysis, fixed prompt order, and absence of teacher implementation data. The study also covers only three models and one prompt realization per task type, and no formal coding-reliability estimate was obtained. Because task type was confounded with turn position within each session, task-level differences should not be interpreted as causal effects of task complexity.

Future work will test the productive-imperfection framework with in-service teachers following IRB approval.

References

Lorin W. Anderson and David R. Krathwohl, editors. 2001. *A Taxonomy for Learning, Teaching, and Assessing: A Revision of Bloom's Taxonomy of Educational Objectives*. Longman.

L. M. Desimone. 2009. Improving impact studies of teachers' professional development: Toward better conceptualizations and measures. *Educational Researcher*, 38(3):181–199.

S. E. Embretson and N. M. Kingston. 2018. Automatic item generation: A more efficient process for developing mathematics achievement items? *Journal of Educational Measurement*, 55(1):112–131.

L. Favero, J. A. Pérez-Ortiz, T. Käser, and N. Oliver. 2025. Leveraging small LLMs for argument mining in education: Argument component identification, classification, and assessment. *arXiv preprint arXiv:2502.14389*.

Gemma Team. 2025. Gemma 3 technical report. *arXiv preprint arXiv:2503.19786*.

M. J. Gierl and T. M. Haladyna, editors. 2012. *Automatic Item Generation: Theory and Practice*. Routledge.

C. S. Große and A. Renkl. 2007. Finding and fixing errors in worked examples: Can this foster learning outcomes? *Learning and Instruction*, 17(6):612–634.

M. J. H. Masum, M. S. H. Majumder, M. M. Rahman, M. Z. Hasan, and F. Al Faisal. 2025. Design and optimization of an offline AI model for enhancing student learning and academic assistance. In *2025 IEEE 7th International Conference on Sustainable Technologies for Industry 5.0 (STI)*, pages 1–6. IEEE.

Microsoft. 2025. Phi-4-mini technical report: Compact yet powerful multimodal language models via mixture-of-loras. *arXiv preprint arXiv:2503.01743*.

W. J. Popham. 2009. Assessment literacy for teachers: Faddish or fundamental? *Theory Into Practice*, 48(1):4–11.

Qwen Team. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.

Z. Reza, A. Mazur, M. T. Dugdale, and R. Ray-Chaudhuri. 2025. Small models, big support: A local LLM framework for educator-centric content creation and assessment with RAG and CAG. *arXiv preprint arXiv:2506.05925*.

F. Wang, Z. Zhang, X. Zhang, Z. Wu, T. Mo, Q. Lu, W. Wang, R. Li, J. Xu, X. Tang, Q. He, Y. Ma, M. Huang, and S. Wang. 2025. A comprehensive survey of small language models in the era of large language models: Techniques, enhancements, applications, collaboration with LLMs, and trustworthiness. *ACM Transactions on Intelligent Systems and Technology*, 16(6):1–87.

A Representative SLM-Generated Item

A.1 Generation Prompt

Prompt. You are writing a Grade 5 multiple-choice math question on adding two fractions with unlike denominators. Make it a multi-step word problem. Provide:

- the question
- four choices A–D (exactly one correct)
- the correct answer
- for each wrong choice, the student misconception it reflects

Use grade-appropriate numbers and wording.

A.2 Representative Phi-4-mini Output

Phi-4-mini, Run 1 — Multi-Step Fraction Addition

Question: Emily has 3/4 of an apple pie slice left over from lunch yesterday; today she eats another third portion as dessert at dinner time. How much does this total fraction represent?

A) 1

B) 7/12

C) More than a whole (2 pieces)

D) Cannot add them together without converting to improper fractions.

Correct Answer: B)

Misconceptions reflected by incorrect answers:

A) Incorrectly summing numerators directly rather than finding common denominators first.

B) Correct understanding of adding 3/4 + 1/3 = (9+4)/12 leading up to total as simplified fraction.

C) Miscalculating overestimation thinking that both fractions contribute beyond a singular whole without the proper addition process considered. Overestimating portions together leads to this misconception.

D) Assuming improper handling of mixed numerical types and not realizing direct methods can still apply using least common denominators for simple fractional sums.

Note. Model output is reproduced verbatim except for typographic formatting.

A.3 Researcher Evaluation

The designated answer is incorrect. The model's own procedure implies

$$\frac{3}{4} + \frac{1}{3} = \frac{9}{12} + \frac{4}{12} = \frac{13}{12},$$

rather than $\frac{7}{12}$. Thus, the response contains a direct inconsistency between the designated key and its own mathematical rationale. The exact value $\frac{13}{12}$ is not offered as an option. Option C correctly indicates that the sum exceeds one, but its parenthetical "2 pieces" is inaccurate; therefore, no listed option is wholly correct. The model nevertheless labels option C as a misconception. This example illustrates how a fluent SLM-generated item can contain multiple interacting defects requiring evaluation beyond answer-key accuracy alone.

B Rubric Grounding and Decision Rules

The six dimensions were designed to translate established principles from assessment literacy, automatic item generation, and cognitive-demand research into teacher-facing evaluation criteria.

D1 Key. Pass if the designated answer is mathematically correct.

D2 Validity. Pass if exactly one option is mathematically correct, the correct answer appears among the options, the key refers unambiguously to one listed option, options are distinct, and options contain response content only.

D3 Alignment. Pass if the item assesses the intended mathematical skill or concept.

D4 Demand. Pass if the reasoning required by the item matches the intended cognitive demand.

D5 Diagnosticity. Pass if all three designated distractors are reachable through identifiable student errors, match the stated misconception, and are supported by mathematically sound rationales.

D6 Clarity. Pass if the item is unambiguous, grade-appropriate, and sufficiently clear for the intended learner.

Principled Approaches to Building AI Partners for Assessing and Supporting Small-Group Collaboration in Classrooms

Peter W. Foltz, Chelsea Chandler, Mon-Lin Monica Ko
Institute of Cognitive Science, University of Colorado Boulder
foltzp@colorado.edu

Abstract

Collaboration is a complex, multifaceted construct, blending cognitive, social, and emotional components; it is context-dependent, dynamic, multimodal and multiparty, properties that resist simple, one-dimensional measurement. We present a principled framework for addressing this complexity that links what is measured, where it is measured, and how evidence is calculated and warranted. This paper (1) synthesizes a taxonomy of key design dimensions for building AI partners that support small-group collaboration; (2) demonstrates how it operationalizes these dimensions in practice. to show how design choices inform how measurement science can be embedded in AI-based collaboration partners.

1 Introduction

As classrooms and workplaces increasingly organize around joint problem-solving rather than individual work, the capacity to collaborate effectively has become one of the most consequential skills a learner can develop. Practicing collaborative problem solving (CPS) helps develop critical thinking, communication, and leadership skills while boosting student engagement (Graesser et al., 2018; Springer et al., 1999; Xu et al., 2023), and employers increasingly rank collaboration alongside problem solving and adaptability as core to career resilience (World Economic Forum, 2025). Unlike many constructs targeted by AI-based assessment, collaboration is a complex, multifaceted construct that blends cognitive, social, and emotional components; it is highly context-dependent; it unfolds dynamically over time rather than as a single snapshot; and it is inherently multimodal and multiparty (Blikstein & Worsley, 2016). These properties make collaboration resistant to one-dimensional measurement and require AI-based approaches capable of modeling nuanced, evolving interpersonal interaction.

Much of the recent focus around AI in education has centered on personalized, one-on-one tutoring, exemplified by large language model (LLM)-based tutors that adapt to an individual learner's needs. Yet learning is not only personal; it is also a social activity situated within a cultural system (Greeno, 1998), and collaborative, open-ended learning environments are critical for fostering deeper engagement and knowledge building (Jeong et al., 2019). Building AI partners that meaningfully support, rather than merely observe, group collaboration therefore requires design principles distinct from those governing single-user tutoring systems.

Evidence-centered design (ECD; Mislevy et al., 2006) provides a principled framework for addressing this complexity by explicitly linking what is measured, where it is measured, and how evidence is calculated and warranted. Applied to collaboration, it clarifies the link between observed discourse or behavior and claims about underlying skills, helping prevent the common failure mode of AI systems optimizing for what is easy to capture (e.g., word counts, click rates) rather than what is construct-relevant (e.g., negotiation, knowledge transfer, perspective-taking). It also supports transparency: by mapping behavioral markers explicitly back to specific evidentiary claims, it reduces the "black-boxiness" of AI-based assessment and supports stakeholder trust among learners, educators, and policymakers.

This paper has three aims. First, we synthesize a taxonomy of key dimensions that guide the principled design of AI partners for collaboration in classrooms (Foltz, 2026). Second, we demonstrate how ECD operationalizes these dimensions in practice. Third, we illustrate these principles by synthesizing work from two real-world implementations developed at the NSF National AI Institute for Student-AI Teaming (iSAT) (D'Mello et al., 2024): the Community

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 350–357
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

Builder (CoBi; Breideband et al., 2023), which supports relationship-building and socially uplifting discourse during open-ended small-group science tasks, and the Jigsaw Interactive Agent (JIA; Ko et al., 2026), which supports productive uncertainty during structured jigsaw collaboration. They illustrate a range of design choices, including AI models, aggregated versus individualized feedback, and varying timing and intervention strategies, that inform how measurement science can be embedded into the next generation of AI-based collaboration partners.

2 Frameworks for Characterizing Collaboration Skills

Assessing collaboration requires identifying the skills to be targeted. Several frameworks have been developed for this purpose, each emphasizing different aspects of cognitive and social functioning. The ATC21S project (Griffin et al., 2012) characterized CPS along participation, perspective-taking, social regulation, and negotiation, focusing on authentic human-to-human interaction but relying on skill descriptors that are difficult to automate. OECD's PISA 2015 Framework (OECD, 2017) defined CPS along a cognitive dimension (task work: exploring, representing, planning, monitoring) crossed with a collaborative dimension (teamwork: establishing shared understanding, taking action, organizing the group), that was scored automatically but within highly scripted, computer-agent interactions. Sun et al. (2020) extended this approach with a focus on joint analysis, negotiation, and knowledge construction, defined via specific, measurable discourse indicators (what is said and done) that support natural language analysis of open, unconstrained conversation.

Despite their differences, these frameworks converge on a common underlying structure: successful collaboration requires both task work (problem-solving strategies, planning, monitoring) and teamwork (communication, coordination, relationship management). Frameworks that link skills to computationally tractable, discourse-level indicators are considerably more amenable to automated assessment than those defined at a broad, human-scored level. What is needed beyond a skills framework, however, is a principled approach for mapping those skills to evidence,

3 Evidence-Centered Design for AI-Based Collaborative Assessment

Evidence-centered design (ECD; Mislevy et al., 2006) provides a comprehensive framework for designing and deploying multimodal, multiparty AI-based assessments by explicitly outlining what is measured (the student model), where it is measured (the task model), and how evidence is calculated (the evidence model). ECD begins with a student model, the claims one wants to make about a learner's or group's knowledge, skills, or abilities, and works backward to define the evidence needed and the tasks that can elicit it. For AI-based assessment, this backward-design logic prevents the common problem of measuring what is easy to capture rather than what is construct-relevant, keeping the analysis focused on construct-relevant features rather than spurious correlations (Foltz, 2026).

Because collaborative data are inherently multimodal, spanning speech, text, gesture, and action logs, ECD requires specifying ahead of time which modalities are relevant in a given context, which features to extract, and how those features support inferences about the target constructs. ECD also enforces explicit reasoning about why a given behavior or pattern of interaction provides evidence for a construct, which removes some of the "black-boxiness" of an AI assessment system: as an AI detects individual markers of behavior, these markers can be explicitly mapped back to specific pieces of evidence, providing a transparent accounting of the analysis process that is essential for stakeholder trust. Prior ECD-grounded implementations illustrate this approach at different scales. For example, the in-task assessment framework (I-TAF; Kerr et al., 2016) applies a CPS ontology to extract behavioral evidence from log data such as mouse clicks and chat messages, while the EPCAL platform (Hao et al., 2017) was built to support control and manipulation of group variables for studying collaboration at scale. As we describe in this paper. we extend this approach with partners to support real-time, open-ended, spoken small-group discourse in authentic K–12 classrooms.

4 A Taxonomy of Design Dimensions for AI-Based Collaborative Support

Identifying collaboration skills and mapping them to evidence is necessary but not sufficient for

building AI technology to assess and support collaboration: these design choices cannot be evaluated in isolation from the practical contexts in which they are applied. Table 1 (adapted from Foltz, 2026), enumerates the key dimensions to consider when designing and deploying AI-based collaborative assessment and support tools, organized into five domains, context and setup, modeling and evidence, timing, feedback and use, and scope and scale, along with each dimension's primary alignment to the ECD student, task, and evidence models, or to practical implementation factors. We summarize each domain below.

4.1 Context and Setup

The context and setup domain addresses the foundational conditions under which an assessment is designed and conducted. The purpose of assessment, whether formative, summative, or diagnostic, carries implications for how the system is structured and what feedback is provided; a formative purpose, in particular, requires mapping identified skills to specific, pedagogically effective feedback. The interaction environment (simulated, live, or hybrid) shapes both the realism and the controllability of the collaboration: live classroom interactions capture authentic dynamics but introduce noise, interruptions, and off-task behavior, whereas simulated settings offer tighter control at the cost of naturalistic behavior. Data-collection choices (spoken audio, video, log files, final products) depend on feasibility and on the richness of evidence required, while also introducing privacy considerations that are heightened when working with minors. Task structure and control range from tightly scripted tasks to open-ended conversation, influencing the range of competencies that can be observed. Finally, teammate embodiment, whether the AI acts as a virtual teammate, provides feedback only, or supports a human-in-the-loop, affects student engagement, trust, and interaction patterns.

4.2 Modeling and Evidence

The modeling and evidence domain concerns what is measured and how evidence is captured, processed, and interpreted. Targeted competencies may include CPS skills, socio-emotional skills, pedagogical moves, or domain-specific knowledge, and can be assessed at varying levels of granularity, including individual, dyad, or whole team. Evidence sources range from verbal and non-verbal discourse markers to group work products, action logs, and observer ratings. Analytic models span a continuum from simple frequency counts to statistical and Bayesian models to fine-tuned language models and LLMs, with more sophisticated multimodal fusion approaches consistently outperforming unimodal counterparts (e.g., Cohn et al., 2024). Communication among team members remains one of the richest sources of evidence about collaborative process and cognition, but spoken discourse must first be processed by automatic speech recognition (ASR), which can exhibit word error rates of 20 to 80% in noisy, real-world classrooms with young speakers, posing a substantial challenge for downstream discourse classification (Southwell et al., 2022).

4.3 Timing and Feedback and Use

The nature and timing of AI feedback can vary widely, from identifying discrete collaborative moves to diagnosing team processes, providing domain-specific scaffolding, or generating incident reports and scores. Data collection and analysis may occur in real time, near-real time, or through post hoc analysis, with each choice trading off immediacy of feedback against scalability and model reliability; a lower-fidelity model that runs in real time may be as effective as a higher-fidelity model requiring post hoc processing. Feedback delivered concurrently with a group task can shape ongoing collaboration but risks disrupting natural interaction, whereas post hoc feedback allows deeper reflection at the cost of losing the opportunity to shape in-the-moment behavior. The target of feedback, whether individual, team, or teacher and evaluator, further shapes its tone, granularity, and actionable quality, and must be balanced against interpretability, specificity, and alignment with broader learning goals.

4.4 Scope and Scale

Scope and scale address broader implementation considerations that determine whether a system can operate effectively and ethically in its intended environment. Scalability of assessment requires planning for pilot phases, small-scale trials, and eventual large-scale deployment, each with distinct infrastructure, training, and quality-assurance needs, as well as organizational buy-in and stakeholder trust. Context sensitivity addresses whether a system is domain-general, domain-

Table 1: Key dimensions for the design of AI-based collaborative assessment and support systems. Adapted from Foltz (2026),

Development Domain	***Collaborative Assessment Dimension***	***Categories***	***Alignment to ECD Layers / Implementation Factors***
Context and setup	Purpose of assessment	Formative, summative, diagnostic	*Implementation*
	Interaction environment	Simulated, live, hybrid	*Task model*
	Data types	Spoken audio, video, log files, student responses	*Implementation*
	Task structure and control	Firmly directed by tasks (multiple choice, SJT), semi-structured, open-ended (e.g., open conversations)	*Task model*
	Teammate embodiment	Virtual teammate, human, human-in-the-loop	*Task model*
Modeling and evidence	Targeted competencies	CPS, social-emotional, pedagogical moves, domain knowledge	*Student model*
	Granularity of analysis	Individual, dyad, whole team	*Student model*
	Collaboration evidence sources	Group work product, verbal expressions, turn-taking, influence, inferred eye gaze, pointing, situation judgment, action logs, survey instruments, observer notes, or scores	*Evidence model*
	Analytic model(s)	Frequency counts, Statistical, Bayesian, semantic feature-based, or LLM-based	*Evidence model*
	Data integration method	Single modality, multimodal, separate, multimodal-integrated	*Evidence model*
	Data access and privacy constraints	Unrestricted, restricted, anonymized	*Implementation*
Timing	Timeliness of data collection and analysis	Real-time, near-real-time, post hoc analysis	*Implementation*
	Timing of AI feedback	Real-time concurrent with group task, near-real-time, post hoc	*Task model*
Feedback and use	Nature of AI feedback	Collaborative moves, team processes, domain scaffolding, reports, critical incidents, scoring	*Task model*
	Target of feedback	Individual, team, teachers, evaluators	*Task model*
Scope and scale	Scalability of assessment	Pilot, small-scale, large-scale	*Implementation*
	Context sensitivity	Domain-general, domain-specific, context-adapted, culturally appropriate	*Implementation*

specific, or context-adapted, and, critically, whether it is culturally appropriate. What constitutes productive participation, appropriate communication, or effective leadership varies substantially across individualistic and collectivistic cultures and different communication norms (Gelfand & Realo, 1999); an AI system trained primarily on Western, industrialized populations that privileges verbal contribution frequency or direct challenge as markers of collaboration quality may systematically disadvantage students from cultures in which such behaviors are read differently. Meaningful accommodation of this variation requires both technical solutions (e.g., diverse training data, domain adaptation, uncertainty quantification) and ongoing participatory collaboration with educators, students, and cultural experts throughout the design and validation process.

5 Illustrative Applications

To ground the taxonomy in Table 1, we describe two AI partners developed at iSAT. It should be noted that the partners are not presented here as competing alternatives; they target different constructs, tasks, and classroom purposes, and were designed to complement each other within a broader ecosystem of collaboration supports. This work reflects a broader commitment to responsible innovation in AI for education, one that shifts the field's attention from personalization toward more explicit emphasis on how AI supports specific forms of learning and how it is designed, both for whom and with whom (e.g., Penuel et al., 2026). Evidence-centered design offers one concrete route to that commitment, by making the constructs, evidence, and feedback choices behind an AI partner explicit rather than incidental. From an ECD-grounded design perspective: each specifies a student model (target constructs), a task model (understanding the classroom activity and data-collection context), and an evidence model (the computational approach linking discourse to constructs). Because each system resolves the dimensions in Table 1 somewhat differently, considering them together lets us illustrate a broader range of the design choices and trade-offs from the taxonomy.

5.1 CoBi: The Community Builder

The Community Builder (CoBi; Breideband et al., 2023) is a multiparty AI partner that targets the relationship dimension of collaboration. Its student model comprises co-negotiated Community Agreements (CAs), namely being respectful, showing commitment to community, and moving and thinking forward, derived from the Sun et al. (2020) CPS framework and grounded in OpenSciEd K–12 science materials aligned with national standards. Its task model situates CoBi within open-ended, small-group science tasks: with teacher guidance, students first co-negotiate concrete examples of each CA; as they then engage in collaborative tasks, CoBi analyzes their spoken discourse in real time; after the task, teachers use CoBi's visualizations to guide collaborative reflection; and, optionally, teachers reveal CoBi's top-ranked "noticings," specific positive exemplars of each CA, to support both collaborative reflection and critical AI literacy.

CoBi's evidence model captures audio via tabletop microphones placed with each student group and analyzes the resulting speech with AI language models trained to recognize evidence of each Community Agreement, outputting a likelihood that a given utterance exemplifies that agreement. Because classroom audio is noisy and automatic speech recognition is error-prone for young speakers, CoBi's models are trained for robustness to transcription error and restrict inference to utterances that can be classified with high confidence (Southwell et al., 2022). Consistent with its evidence and task models, results are deliberately aggregated at the classroom level rather than reported to individuals, reducing distraction and surveillance while preserving CoBi's formative, reflective purpose. Feedback is delivered through real-time visualizations, including a growing "tree" of noticings, displayed to the whole class.

CoBi is built on a scalable, modular cloud architecture and has been deployed with hundreds of middle-school student groups across multiple school districts, requiring only a Chromebook and an affordable USB microphone per group. Preliminary classroom testing found that students responded positively to CoBi's reflective visualizations, though when confronted with their own results, students tended to attribute shortcomings to their own collaboration behavior rather than to possible errors in CoBi's classification, an important consideration for how transparency and AI-literacy activities are designed (Breideband et al., 2023).

5.2 JIA: The Jigsaw Interactive Agent

The Jigsaw Interactive Agent (JIA; Ko et al., 2026) targets a different construct and a more structured task. In a jigsaw activity, students first develop individual expertise on different aspects of a topic and then work together in a small group to pool that expertise and solve a problem that requires everyone's contribution (Aronson, 1978). Its student model comprises four discourse-based intervention types derived from research on teacher facilitation of productive uncertainty, namely support, connect, problematize, and stabilize (Watkins & Manz, 2022), reflecting a pedagogical stance that seeks to help students surface and productively work through, rather than resolve away, the non-obvious and contingent aspects of their collective sensemaking (Manz & Suárez, 2018). Its task model situates JIA within a jigsaw activity embedded in a co-designed, three-week AI-literacy curriculum unit, in which students first develop individual expertise on one of five content-moderation approaches and then return to their small group to pool that expertise and rank the approaches by effectiveness, a task that structurally requires collaboration, since no single student holds complete information.

JIA's evidence model was initially deliberately varied to compare three dialog-policy approaches, each analyzing the immediately preceding minute of group discourse and triggering one of the four intervention types directly into the group's chat interface: a rule-based approach that selected interventions from indicators of participation structure, task engagement, and collaboration quality; an LLM-based approach (Llama 70B) that operated directly on ASR-derived transcripts, summarizing the discussion and justifying its chosen intervention in natural language; and a hybrid approach that supplied the same LLM with both the conversational transcript and the rule-based indicators. All three approaches contended with substantial automatic speech recognition error, approximately 58% word error rate in benchmark classroom conditions.

JIA was formatively evaluated with 64 small groups of 8th-grade students across 11 classrooms in two U.S. school districts, of which 16 groups with moderate-to-high engagement were the focus

of detailed human-annotated analysis (Ko et al., 2026). The LLM- and hybrid-based approaches were significantly more likely to generate connect interventions and were judged more appropriate by human coders (82% and 87%, respectively) than the rule-based approach (63%), which defaulted to more repetitive support interventions. Students reacted more often to the LLM approach (49%) than to the rule-based (24%) or hybrid (29%) approaches and displayed a wider range of emotional responses, while connect and problematize interventions elicited more elaborated uptake than support interventions..

6 The Taxonomy's Design Space

CoBi and JIA illustrate how the dimensions in Table 1 can be resolved differently depending on a system's construct, task, and purpose. In the modeling and evidence domain, JIA's formative comparison of rule-based, LLM-based, and hybrid dialog policies shows one way to weigh interpretability and computational cost against contextual responsiveness. Its LLM- and hybrid-based approaches produced more contextually responsive, construct-aligned feedback, more often connecting ideas across group members rather than defaulting to generic encouragement, but also elicited more variable emotional reactions from students, including occasional dismissiveness alongside enthusiasm. This illustrates a trade-off inherent to the modeling and evidence domain: more flexible, generative approaches can better track construct-relevant nuance in open discourse, at some cost to predictability.

The systems also illustrate the range of choices in the feedback and use domain. CoBi's classroom-level aggregation reflects a design decision aligned with its evidence and task models: because CoBi is a reflection-support tool oriented toward the relationship dimension of collaboration, aggregating feedback across groups reduces distraction and the risk that students experience the tool as surveillance, while preserving its formative, reflective function. JIA instead delivers feedback directly to each small group in real time, functioning as an embedded teammate whose goal, supporting in-the-moment productive uncertainty, calls for timely, group-specific intervention rather than delayed, aggregated reflection. Together, the two systems show how the appropriate target of feedback follows from the construct and purpose being assessed rather than from a single best practice.

Timing choices follow a similar logic. CoBi's visualizations update continuously but are intended primarily for use during teacher-led reflection after a task, while JIA intervenes proactively at fixed intervals during group work itself. Each choice reflects the timing and feedback domain's central trade-off between disruption and immediacy: JIA's interval-based design was intended to support students in the moment without overwhelming them, while CoBi's post-task reflection avoids interrupting collaboration but limits the chance to shape it as it happens.

Finally, the two systems illustrate different points along the scope and scale domain. CoBi has been deployed at substantial scale, across hundreds of groups and multiple school districts, on a production cloud architecture. JIA remains at the formative pilot stage, with 64 groups and detailed analysis of 16, reflecting the additional validation an actively intervening agent typically requires before wider classroom use. Despite these differences in maturity, both systems demonstrate that ECD-grounded design enables defensible, real-time analysis of noisy, multiparty classroom discourse, in part by restricting inference to high-confidence evidence and by validating models against human-annotated criterion data (Foltz, 2026).

7 Toward Validated, Ethical, and Scalable AI Collaboration Partners

The AI collaborative partners demonstrate that principled, grounded design choices, including what is measured, how evidence is extracted from imperfect multimodal data, and how and to whom feedback is delivered, materially shape both the validity and the classroom impact of AI collaboration partners. Implementing AI partners at scale also requires confronting challenges that fall outside the immediate scope of Table 1 but are inseparable from responsible implementation. Multimodal data collection, particularly of children's speech, and potentially video, gesture, or eye gaze, raises privacy and consent challenges that a purely technical framework cannot resolve on its own; informed consent must move from a binary opt-in/opt-out decision toward a continuous, measurable process that maintains transparency and accountability throughout data collection (Penuel et al., 2026). Regulatory frameworks such

as FERPA, COPPA, and PPRA impose overlapping but distinct requirements that researchers and developers must navigate to protect students' rights and dignity. Demonstrating the educational value that such tools provide, for instance through transparent, critiquable feedback in the spirit of CoBi's noticings, is likely a prerequisite for broader stakeholder buy-in.

A second, related challenge is generalizability across curricula, cultural contexts, and task structures. As emphasized in the scope and scale domain, collaboration norms are not culturally universal, and systems trained on narrow populations risk misreading, or systematically disadvantaging, students whose collaborative styles differ from those of the training population. Addressing this will require diverse training data, domain adaptation and uncertainty-aware modeling, and, above all, sustained participatory design with educators, students, and cultural experts from the earliest stages of competency modeling through validation.

Collaboration is difficult to measure precisely because it is not one thing: it is cognitive and social, momentary and cumulative, individual and collective. Evidence-centered design offers a way to hold that complexity accountable to explicit claims, tasks, and evidence rather than to whatever happens to be easiest to compute. Applying ECD to real-world AI partners, shows both the promise and the trade-offs of principled design: aggregated versus individualized feedback, different points along the spectrum from validated classifiers to more flexible but variable generative models, and reflective versus in-the-moment intervention.

As AI-based collaboration partners move from research pilots toward broader classroom use, continued attention to transparent, evidence-based, and human-centered design, anchored in frameworks, will be essential to ensuring that these systems serve, rather than undermine, the students and teachers they are meant to support. The present work contributes a generalizable taxonomy of design dimensions, grounded in measurement science, for researchers and developers building the next generation of AI-based collaborative assessment and support tools. By making explicit the choices bundled into any AI collaboration partner, choices too often left implicit, we hope to give the measurement community a shared vocabulary for evaluating, comparing, and improving these systems as they move from classroom pilots toward broader, validated deployment.

Acknowledgments

This research was supported by the NSF National AI Institute for Student-AI Teaming (iSAT) under grants DRL 2019805 and DRL 2454151. The opinions expressed are those of the authors and do not represent the views of the NSF.

References

Aronson, E. (1978). *The jigsaw classroom*. Sage.

Blikstein, P., & Worsley, M. (2016). Multimodal Learning Analytics: A methodological framework for research in constructivist learning. *Journal of Learning Analytics,* 3(2), 220–238.

Breideband, T., Bush, J., Chandler, C., Chang, M., Dickler, R., Foltz, P., … & D'Mello, S. (2023, October). The Community Builder (CoBi): Helping students to develop better small group collaborative learning skills. In Companion Publication of the 2023 *Conference on Computer Supported Cooperative Work and Social Computing* (pp. 376–380). ACM. https://doi.org/10.1145/3584931.3607498

Cohn, C., Davalos, E., Vatral, C., Fonteles, J. H., Wang, H. D., Ma, M., & Biswas, G. (2024). Multimodal methods for analyzing learning and training environments: A systematic literature review. *arXiv preprint* arXiv:2408.14491.

D'Mello S. K., Biddy Q., Breideband T., Bush J. B., Chang M. A., Cortez A., Flanigan J., Foltz P. W., Gorman J. C., Hirshfield L. M., Ko M.-L. M., Krishnaswamy N., Lieber R., Martin J., Palmer M. S., Penuel W. R., Walker L., Whitehill J. (2024). From learning optimization to learner flourishing: Reimagining AI in education at the Institute for Student-AI Teaming (iSAT). *AI Magazine*.

Foltz, P. W. (2026). Key considerations for assessing and supporting collaboration with AI. In (X. Xiong, M. D. Shermis & J. Xiong, Eds.). *The Role of AI in Assessment* (pp. 287–311). Routledge, London. https://doi.org/10.4324/9781003621522-12

Gelfand, M. J., & Realo, A. (1999). Individualism-collectivism and accountability in intergroup negotiations. *Journal of Applied Psychology*, 84(5), 721–736.

Graesser, A. C., Fiore, S. M., Greiff, S., Andrews-Todd, J., Foltz, P. W., & Hesse, F. W. (2018). Advancing the science of collaborative problem solving. *Psychological Science in the Public Interest*, 19(2), 59–92. https://doi.org/10.1177/1529100618808244

Greeno, J. G. (1998). The situativity of knowing, learning, and research. *American Psychologist*, 53(1), 5–26.

Griffin, P., McGaw, B., & Care, E. (Eds.). (2012). *Assessment and teaching of 21st century skills*. Springer.

Hao, J., Liu, L., von Davier, A. A., Lederer, N., Zapata-Rivera, D., Jakl, P., & Bakkenson, M. (2017). EPCAL: ETS platform for collaborative assessment and learning. *ETS Research Report Series*, 2017(1), 1–14.

Jeong, H., Hmelo-Silver, C. E., & Jo, K. (2019). Ten years of Computer-Supported Collaborative Learning: A meta-analysis of CSCL in STEM education during 2005–2014. *Educational Research Review,* 28, 100284. https://doi.org/10.1016/j.edurev.2019.100284

Kerr, D., Andrews, J. J., & Mislevy, R. J. (2016). The in-task assessment framework for behavioral data. In A. A. Rupp & J. P. Leighton (Eds.), The Wiley *Handbook of Cognition and Assessment: Frameworks, Methodologies, and Applications* (pp. 472–507). Wiley.

Ko, M. M., Chandler, C., Rose, S., Cline, B., Watts, E., Reitman, J., Foltz, P. W., & D'Mello, S. K. (2016). AI Partners that Support Productive Uncertainty Within "Jigsaw" Activities During Small Group Collaborative Learning in Classrooms. In *International Conference on Artificial Intelligence in Education* (pp. 608-622). Cham: Springer Nature Switzerland.

Manz, E., & Suárez, E. (2018). Supporting teachers to negotiate uncertainty for science, students, and teaching. *Science Education*, 102(4), 771–795.

Mislevy, R. J., Steinberg, L. S., Almond, R. G., & Lukas, J. F. (2006). Concepts, terminology, and basic models of evidence-centered design. In D. M. Williamson, R. J. Mislevy, & I. I. Bejar (Eds.), *Automated Scoring of Complex Tasks in Computer-Based Testing* (pp. 15–47). Routledge.

OECD. (2017). *PISA 2015 results: Collaborative problem solving*. OECD Publishing.

Penuel, W. R., Philip, T. M., A. Chang, M., Sumner, T., & D'Mello, S. K. (2026). Responsible Innovation in Designing AI for Education: Attending to the "How," the "for What," the "for Whom," and the "With Whom" in a Rapidly Growing Field. *Educational Researcher*, 55(2), 134-141. https://doi.org/10.3102/0013189X251389879

Southwell, R., Pugh, S., Perkoff, M., Clevenger, C., Bush, J. B., Lieber, R., … & D'Mello, S. K. (2022). Challenges and feasibility of automatic speech recognition for modeling collaborative discourse in classrooms. *Proceedings of the 15th International Educational Data Mining Conference (EDM 2022)* (pp. 302–315).

Springer, L., Stanne, M. E., & Donovan, S. S. (1999). Effects of small group learning on undergraduates in science, mathematics, engineering, and technology: A meta-analysis. *Review of Educational Research*, 69(1), 21–51.

Sun, C., Shute, V., Stewart, A., Yonehiro, J., Duran, N., & D'Mello, S. K. (2020). Toward a generalized competency model of collaborative problem solving. *Computers & Education,* 143, 103672.

Watkins, J., & Manz, E. (2022). Characterizing pedagogical decision points in sense-making conversations motivated by scientific uncertainty. *Science Education,* 106(6), 1408–1441.

World Economic Forum. (2025). *The future of jobs report 2025.* https://www.weforum.org/publications/the-future-of-jobs-report-2025/in-full/3-skills-outlook/

Xu, E., Wang, W., & Wang, Q. (2023). The effectiveness of collaborative problem solving in promoting students' critical thinking: A meta-analysis based on empirical literature. *Humanities and Social Sciences Communications, 10*(1), 1–11.

Measuring Writing Proficiency Through Keystroke Process Data: A Cross-School Generalizability Study

Damilola Babalola[1] **Naga Buddarapu**[1] **Caleb Scott**[1] **Piotr Mitros**[2] **Paul Deane**[2] **Collin Lynch**[1]
[1] North Carolina State University
[2] Educational Testing Service
{djbabalo,vbuddar,cwscott3,cflynch}@ncsu.edu
{pmitros,pdeane}@ets.org

Abstract

Keystroke logging offers a scalable way to observe writing as it unfolds, but its value for educational measurement depends on whether process features are valid and stable across schools, grades, and writing contexts. We analyze 8,576 authentic classroom documents from 423 seventh- and eighth-grade students in two U.S. middle schools. Twenty-four theory-motivated keystroke measures were organized into fluency, revision, complexity, and engagement families and expanded through distributional summaries of timing measures. We examine associations between process features and incoming state English Language Arts (ELA) proficiency, the interpretability and stability of latent factor structures, and the consistency of gender-linked differences across schools and writing contexts. Results show that several process features, particularly those related to engagement and timing, were associated with ELA proficiency, with the strongest and most consistent patterns emerging in Grade 8. These associations showed moderate replication across schools and writing contexts. Factor analyses identified recurring dimensions related to engagement, revision, timing variability, and planning, although their structure was not fully stable across cohorts. Gender-linked differences were also inconsistent across grades and schools. Overall, the findings support keystroke measures as promising process signals for AI-enabled formative assessment while highlighting the need for local validation, adequate sampling, and fairness checks before deployment.

1 Introduction

Writing proficiency is expressed not only in the quality of a finished text but also in the processes through which that text is planned, produced, reviewed, and revised. Cognitive models describe writing as a recursive, self-regulated activity in which writers allocate limited attention and working-memory resources across planning, translating, and reviewing (Flower and Hayes, 1981; Graham et al., 2017; Zimmerman and Risemberg, 1997). Traditional product-focused assessments observe the outcome of these processes, emphasizing text quality, coherence, and surface-level correctness, but provide limited evidence about how students reached it. With advances in educational data mining, keystroke logging now enables the fine-grained observation of writing in real-time, capturing insertions, deletions, pauses, and navigational behaviors (Leijten and Van Waes, 2013; Baaijen et al., 2012).

This process evidence is increasingly relevant to artificial intelligence (AI) in educational measurement. Automated writing evaluation has historically relied primarily on features extracted from the final text, but more recent studies have used keystroke traces in machine-learning models to predict writing quality, identify students who may need support, distinguish human from AI-generated text, and differentiate planning strategies across student clusters (Conijn et al., 2022; Litman et al., 2022; Jiang et al., 2024). More recent measurement research has also modeled raw keystroke sequences with large language model architectures to infer stable characteristics of writers (Zhang et al., 2025). These applications make keystroke data attractive as an additional behavioral channel for AI-enabled assessment, but they also raise a basic measurement question: before process measures are used for prediction, feedback, or automated decision-making, are the signals sufficiently stable and generalizable across educational contexts?

Existing keystroke studies provide evidence that timing, fluency, revision, and time-on-task measures can relate to writing quality, but much of that evidence comes from controlled or time-limited tasks, adult or postsecondary writers, or a single assessment context (Leijten and Van Waes, 2013;

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 358–368
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

Conijn et al., 2022, 2019). Authentic middle-school writing is less standardized: students write across subjects, assignments, teachers, and time spans, and the meaning of a pause or revision may vary with task demands (Conijn et al., 2019). A signal that is predictive in one cohort may therefore fail to transfer to another. This is especially important for classroom-embedded AI systems, where instability can be mistaken for student differences and where demographic patterns can become embedded in automated interpretations.

We address this issue using a full academic year of naturally occurring writing from two U.S. middle schools. The study treats cross-context replication as a measurement requirement rather than as an afterthought. RQ1, our primary focus, asks which keystroke-derived behaviors are associated with students' incoming ELA proficiency and whether those associations replicate across schools, grades, and writing contexts. RQ2 asks whether the features form interpretable latent dimensions and whether those structures generalize. RQ3 provides an exploratory fairness-oriented check of whether gender-linked feature differences are stable or context-bound. Together, the analyses evaluate whether process features are plausible candidates for classroom-embedded, AI-supported formative measurement, where process evidence may complement established approaches to formative writing assessment (Graham et al., 2015).

2 Background

2.1 From Keystrokes to Writing Constructs

Keystroke logging has become instrumental in identifying writing strategies and understanding cognitive processes by making the process underlying the written product observable (Deane et al., 2026). To interpret keystroke data meaningfully, low-level behavioral metrics must be linked to higher-order cognitive constructs (Deane et al., 2008; Odendahl and Deane, 2018). The present study organizes its features into four theoretical families grounded in established models of writing cognition:

- **Fluency** reflects the efficiency and rhythm with which ideas are translated into text, including insertion activity and inter-keystroke intervals within and between words (Van Waes and Leijten, 2015).
- **Revision** captures self-monitoring and correction through edits, backspaces, and deletion sequences, consistent with models of writing as a self-regulated process (Baaijen et al., 2012; Graham et al., 2017; Zimmerman and Risemberg, 1997).
- **Complexity** captures timing and behavior at clause and punctuation boundaries, where sentence-level and rhetorical planning may become visible (Flower and Hayes, 1981; Deane et al., 2008).
- **Engagement** captures sustained allocation of attention through active writing, reading, navigation, and menu time (Conijn et al., 2019).

The feature families are theoretical reference frames rather than claims that a single keystroke event uniquely identifies a cognitive state. Pauses at different textual locations can reflect different underlying processes, including lexical processing and higher-level planning, making one-to-one mappings between pause behavior and cognition difficult (Conijn et al., 2019; Révész et al., 2026). This approach also motivates the latent-structure analysis in RQ2: if multiple observed features repeatedly cluster in interpretable ways, that convergence provides additional evidence about construct representation, while instability cautions against universal feature-to-construct mappings.

2.2 Keystroke Data in AI-Enabled Assessment

Process data are already being incorporated into AI and machine-learning approaches to writing assessment. Conijn et al. (2022) trained regression and classification models on keystroke features to test early prediction of writing quality and identification of writers who might need support. Jiang et al. (2024) used keystroke behavior patterns in machine-learning detectors of nonauthentic writing in a large-scale assessment and explicitly evaluated fairness across demographic groups. Zhang et al. (2025) described both theory-driven feature modeling and an approach in which raw keystroke logs were input to large language models to infer stable writer characteristics. These studies show that keystroke traces are not merely descriptive research data; they are becoming inputs to AI-supported measurement systems.

More broadly, unobtrusively collected behavioral traces have also been used for "stealth" assessment, in which learner performance is in-

ferred from activity generated during authentic tasks rather than from a separate test (Shute et al., 2016). That trajectory strengthens the need for validation studies that test portability before deployment. Keystroke features can be sensitive to task demands (Conijn et al., 2019), meaning that a signal that is statistically useful in one dataset may not have the same distribution or interpretation elsewhere. Our emphasis on cross-school, cross-grade, and cross-context stability therefore complements predictive-model research: it asks whether the behavioral signals themselves behave consistently enough to support downstream AI use.

3 Data and Methods

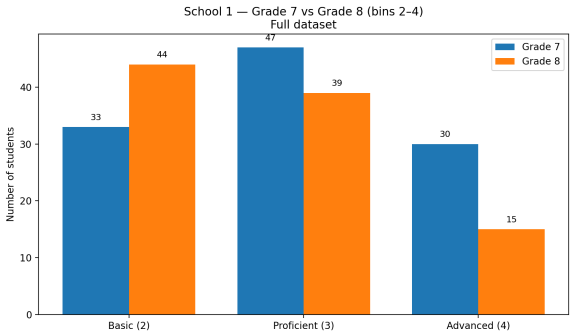
3.1 Participants, Writing Data, and Proficiency Measure

Data were collected during the 2023-2024 academic year from students in Grades 7 and 8 at two U.S. middle schools participating in a research project. School 1 is a large urban school in Utah (857 students; 42% White, 42% Hispanic, 8% multiracial, 3.5% African American, 3% Asian), and School 2 is a suburban school in Pennsylvania (713 students; 75% White, 7% Hispanic, 7% Asian American, 4.2% African American). A classroom-embedded Chrome extension captured Google Docs writing activity, including insertions, deletions, pauses, navigation, and reading-related events. The logs span ELA assignments (e.g., essays, narratives, text-dependent analyses) and non-ELA writing (e.g., lab reports, notes). The non-ELA category is important because students engage in substantial writing activities across subjects, while explicit writing instruction typically occurs in ELA. The study was conducted under Institutional Review Board approval with pseudonymized data handling.

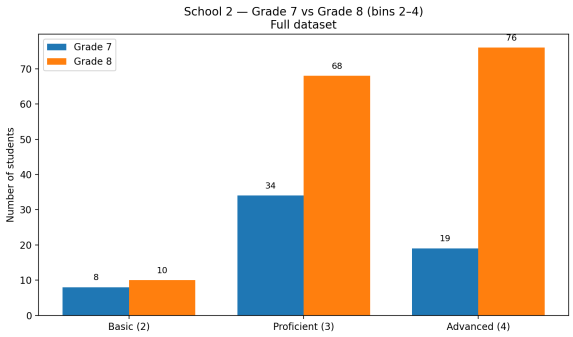
The initial collection contained 9,831 documents from 544 students. Documents shorter than 50 words were excluded. Students' prior-year (2022–2023) state ELA proficiency scores served as the reference performance measure. Although assessments differ by state, scores were harmonized into four bins (1 = Below Basic, 2 = Basic, 3 = Proficient, 4 = Advanced). Because School 2 contained no Below-Basic cases and School 1 contained few, bin 1 was excluded, and the remaining three levels were used to ensure comparability across schools. The final corpus consisted of 8,576 documents (School 1: 3,641; School 2: 4,935) from 423 students. We analyze two writing contexts. The Full dataset contains all eligible classroom writing while the ELA subset focuses on extended ELA writing identified through assignment-title keywords (e.g., Essay, Narrative, Argument, Analysis, TDA) and minimum thresholds of 200 words and 50 insertion events. Comparing the two contexts tests whether associations observed in ELA writing also hold across the broader scope of classroom writing. Student demographics (gender, grade, class period, English Learner status) were obtained from school rosters. Table 1 summarizes sample sizes by school, grade, and dataset. Figures 1a and 1b show the distribution of both schools for the full dataset.

		Students		Documents	
School	**Context**	**G7**	**G8**	**G7**	**G8**
School 1	Full	110	98	1,840	1,801
School 1	ELA	102	93	510	417
School 2	Full	61	154	989	3,946
School 2	ELA	55	153	170	806

Table 1: Analytic sample by school, context, grade, and document count.



(a) School 1



(b) School 2

Figure 1: Distribution of state ELA proficiency levels by grade and school.

Dimension	Keystroke measures
Fluency	Number of inserts; in-word inserts; between-word inserts; word-start inserts; IKI between inserts (interval between successive insertions); IKI within-word (interval between letters); IKI between words (interval between words); IKI before-word (interval preceding word onset); IKI at word-start (interval at word initiation).
Revision	Before-word inserts; total edits; starting backspaces (at word start); continuing backspaces (sequential deletions across words); IKI starting backspace (interval before deletion onset); IKI continuing backspace (interval during deletion sequences); edit duration.
Complexity	Clause-break inserts (marking syntactic boundaries); minor punctuation inserts (within-sentence punctuation); IKI at clause-break (interval at clause boundaries); IKI at minor punctuation (interval surrounding within-sentence punctuation).
Engagement	Overall menu time; overall reading time (inactive but focused); overall navigation time (cursor, scroll, and arrow-key movement); overall active writing time.

Table 2: Theory-motivated keystroke feature families used in the analysis.

3.2 Feature Extraction

The keystroke logs were transformed into quantitative indicators of writing behavior from timestamped keypress, deletion, pause, and navigation events. Following the framework in Section 2.1, the 24 extracted measures were grouped into fluency, revision, complexity, and engagement dimensions (Table 2). Ten measures were frequency based and fourteen were time based. Inter-keystroke interval (IKI) denotes the elapsed time between successive keystroke or editing events and captures temporal aspects of writing behavior.

To reduce the influence of heavy-tailed timing distributions, time-based measures were transformed using $\log(1 + x)$, and pauses exceeding 120 seconds were truncated as likely off-task intervals. For timing measures other than overall writing, reading, navigation, and menu time, we calculated the mean, median, standard deviation, and variance, yielding 54 analytic variables (Leijten and Van Waes, 2013; Baaijen et al., 2012; Conijn et al., 2022, 2019; Van Waes and Leijten, 2015). In the variable names for the analyses, *mean*, *median*, *std*, and *var* denote the distributional summaries, while *n* denotes a count.

3.3 Analyses

- **RQ1:** Within each School $\times$ Grade $\times$ Context cohort, we computed Spearman rank correlations between each feature and the ordered incoming ELA proficiency level. Benjamini-Hochberg false discovery rate (FDR) correction (Benjamini and Hochberg, 1995) was applied at $\alpha = .10$. To test generalizability, sets of significant features were compared using Jaccard similarity (Jaccard, 1901) across (a) ELA versus Full contexts within school and grade, (b) Grade 7 versus Grade 8 within school and context, and (c) School 1 vs School 2 within grade and context. Directional agreement was calculated for shared features, and correlation analyses used students with valid proficiency scores and available feature data; consequently, analytic N may be slightly smaller than the cohort counts reported in Table 1.

- **RQ2:** Exploratory factor analysis (EFA) was conducted separately by cohort using standardized features, principal-axis factoring, and Promax rotation. Extreme values were winsorized at the 1% tails. Cohorts required $N \geq 40$, and sampling adequacy was assessed using the Kaiser-Meyer-Olkin (KMO) statistic (Kaiser, 1960), with values below .60 treated as marginal. Factorability was also assessed using Bartlett's test of sphericity (Bartlett, 1950), and loadings with $|\lambda| \geq .30$ were treated as salient. Factor retention used parallel analysis (Horn, 1965) together with the eigenvalue-greater-than-one criterion, conservatively retaining the smaller estimate. Tucker's congruence coefficient (Lorenzo-Seva and ten Berge, 2006) quantified similarity between factor-loading structures across contexts, grades, and schools. Because several cohorts are modest in size relative to the number of variables, low congruence is interpreted as potentially reflecting both genuine context sensitivity and estimation instability (MacCallum et al., 1999).

- **RQ3:** Gender comparisons were exploratory and limited to binary roster labels (Female/Male), the only consistently available demographic field across sites. Two-sided Mann-Whitney U tests were applied within cohorts with FDR correction at $\alpha = .10$,

and Cliff's delta (δ) summarized effect magnitude and direction (Cliff, 1993). Analyses required at least 10 students per gender group and 30 students per cohort. Jaccard overlap and directional consistency assessed replication across contexts, grades, and schools.

4 Results

4.1 RQ1: Proficiency Associations and Generalizability

The clearest proficiency associations emerged in Grade 8 (detailed feature-level results are reported in Appendix A.1). In School 1 ELA writing, overall reading time ($\rho = .35$), between-word timing variability ($\rho = .34$), and clause-break timing variability ($\rho = .30$) were positively associated with incoming ELA proficiency, while longer mean clause-boundary pauses were negatively associated ($\rho \approx -.24$). The Full dataset showed a similar pattern, with reading, active-writing, and between-word activity among the strongest signals.

School 2 showed an even clearer Grade 8 profile: navigation, reading, and active-writing time correlated positively with proficiency ($\rho \approx .35$–$.41$), whereas mean and median clause-boundary timing correlated negatively ($\rho \approx -.40$). Together, these results suggest that students with higher incoming ELA proficiency were characterized not simply by greater text production, but by sustained engagement and distinct temporal patterns, including shorter pauses at linguistic boundaries.

Neither school's Grade 7 ELA subset produced significant feature–proficiency associations after FDR correction. In the Full datasets, School 1 Grade 7 showed three associations involving active writing, menu time, and between-word insertions, whereas School 2 Grade 7 showed only a navigation-time association. Differences in cohort size, score distributions, and range restriction prevent interpreting this pattern as developmental change.

Replication analyses nevertheless reinforced the Grade 8 result (Table 3). ELA–Full overlap was moderate in both schools ($J = .53$ and $.58$) with perfect directional agreement, while cross-school Grade 8 overlap reached $J = .36$ in ELA and $J = .58$ in the Full dataset. Thus, the strongest Grade 8 fluency and engagement signals generalized beyond a single school or writing context.

Comparison	Cohort	Jacc.	Overlap	Sign
ELA vs. Full	S1, G8	0.53	8	1.00
ELA vs. Full	S2, G8	0.58	15	1.00
S1 vs. S2	ELA, G8	0.36	9	0.89
S1 vs. S2	Full, G8	0.58	5	1.00
G7 vs. G8	S1, Full	0.27	3	1.00
G7 vs. G8	S2, Full	0.05	1	1.00

Table 3: Stability of significant RQ1 feature–proficiency associations. Overlap is the number of shared significant features; Sign is the proportion with matching direction.

4.2 RQ2: Latent Behavioral Structure

Exploratory factor analysis yielded five to eight factors across cohorts. Despite variation in exact composition, several recurring behavioral dimensions emerged. A global-engagement factor commonly combined active-writing, reading, and navigation time. Revision-related factors separated rapid local revision from more sustained editing, with insertion and backspace intervals reflecting revision fluency and longer edit durations capturing deliberate multi-step refinement. Timing variability formed another recurring dimension, while pre-word pauses and clause- or punctuation-level timing reflected planning and linguistic rhythm. These planning and clause-level patterns were more prominent in Grade 8.

The Grade 8 ELA solutions illustrate both the recurring behavioral dimensions and differences in their precise composition across schools. In both schools, active-writing, navigation, reading, and menu time clustered within a global-engagement factor, while continuing-backspace and edit-duration measures formed a sustained-revision factor. The ordering and composition of the remaining factors differed across schools (Appendix A.2).

Bartlett's tests of sphericity were significant across all analyzed cohorts, supporting factorability. Sampling adequacy ranged from marginal to acceptable (KMO = .52–.78), with the weakest values occurring in School 2 Grade 7. Although these recurring dimensions were interpretable, their exact factor composition was not stable across cohorts. Cross-context congruence was negligible in Grade 7 but higher in Grade 8, particularly for School 1 ($\phi = .66$). Cross-grade and cross-school congruence remained below conventional similarity thresholds.

Thus, the data support recurring families of en-

School	Grade	ELA KMO	Full KMO	Congr.
School 1	G7	.782 (6)	.757 (8)	.01
	G8	.767 (6)	.764 (6)	.66
School 2	G7	.520 (5)	.609 (5)	.07
	G8	.725 (6)	.721 (7)	.37

Table 4: RQ2 sampling adequacy and cross-context factor stability. KMO values are followed in parentheses by the number of retained factors. Congr. denotes ELA–Full Tucker congruence.

Comparison	Cohort	Jacc.	Overlap	Interpretation
ELA vs. Full	S1, G7	.57	12	Strong within-cohort replication
ELA vs. Full	S2, G8	.50	6	Moderate within-cohort replication
G7 vs. G8	All pairs	.00	0	No cross-grade persistence
S1 vs. S2	Full, G7	.11	2	Minimal cross-school replication

Table 5: RQ3 stability of significant gender-linked feature differences. S1 and S2 denote Schools 1 and 2; G7 and G8 denote Grades 7 and 8.

gagement, revision, timing variability, and planning behavior, but not a single invariant latent structure. Given the modest size of several cohorts, this instability may reflect both genuine context sensitivity and factor-estimation uncertainty (MacCallum et al., 1999).

4.3 RQ3: Exploratory Gender Differences

Gender-linked differences were less stable than proficiency associations. In School 1 Grade 7, female students showed greater between-word insertion activity ($\delta = -.57$), navigation ($\delta = -.52$), active-writing activity ($\delta = -.51$), and reading activity ($\delta = -.48$), whereas male students showed greater clause-boundary timing variability ($\delta = .52$). These effects replicated across ELA and Full writing contexts. In contrast, no significant gender differences were observed for School 1 Grade 8.

School 2 showed a different pattern. Grade 7 ELA produced no significant differences, while the Full dataset showed only a small number of male-dominant word-start timing effects. In Grade 8, female students showed somewhat greater engagement, whereas male students showed longer or more variable clause- and word-start pauses. Detailed feature-level gender comparisons are reported in Appendix A.3.

Within-cohort replication could therefore be moderate to strong, but cross-grade and cross-school persistence was minimal or absent (Table 5). The results do not support gender as a stable explanation of keystroke behavior; instead, observed differences appear contingent on cohort and writing context.

5 Discussion

This study examined whether keystroke-derived writing-process measures remain meaningful across schools, grades, and writing contexts. The strongest evidence emerged in Grade 8, where higher incoming ELA proficiency was consistently associated with greater reading, navigation, and active-writing activity and with shorter pauses at clause boundaries. Replication across schools and across ELA and broader classroom writing suggests that these relationships are not confined to a single institutional or task context. These findings are broadly consistent with prior work showing that keystroke-derived timing, fluency, revision, and engagement measures contain information related to writing performance (Conijn et al., 2022; Van Waes and Leijten, 2015). At the same time, prior evidence that keystroke features are sensitive to writing-task demands (Conijn et al., 2019) makes the cross-context replication observed here particularly important.

At the same time, the findings caution against treating all process measures as portable indicators. Grade 7 associations were sparse, factor structures were not invariant, and gender-linked differences did not replicate consistently. These results matter for AI-enabled educational measurement because predictive usefulness within one cohort does not by itself establish the validity of an interpretation or its transferability to other populations and contexts (American Educational Research Association et al., 2014). A model may exploit local timing, task, or demographic patterns while performing poorly or differently elsewhere.

The results therefore support a staged approach to using keystroke data in AI systems. Candidate features should first be theoretically motivated and locally validated, then tested for replication across relevant contexts and subgroups. Predictive models should be evaluated not only for accuracy but also for stability and subgroup performance, particularly because automated scoring models can reproduce measurement and algorithmic biases through their features and modeling decisions (Johnson et al., 2022; Jiang et al., 2024).

Process-based feedback should likewise remain tied to observable behavior rather than treating isolated keystroke events as direct evidence of internal cognitive states.

The recurring fluency, revision, engagement, and planning-related patterns provide some support for the interpretability of the feature design, but the unstable factor structures suggest that these categories are better treated as organizing frameworks than fixed latent scales. This interpretation is consistent with prior cautions that observable keystroke behaviors do not map uniquely onto higher-order writing processes (Deane et al., 2008; Conijn et al., 2019). For downstream AI applications, keystroke features may therefore be most useful as context-sensitive evidence sources rather than universal proxies for proficiency.

Among the features examined, reading, navigation, active-writing time, and clause-boundary timing appear most promising because their relationships with proficiency replicated across Grade 8 schools and writing contexts. By contrast, features or factor structures observed in only one cohort should be treated as locally informative until they replicate elsewhere. This distinction is important as keystroke data are increasingly used for writing-quality prediction, detection of nonauthentic text, and other AI-supported assessment tasks (Conijn et al., 2022; Jiang et al., 2024; Zhang et al., 2025; Deane et al., 2026).

6 Limitations and Future Work

First, excluding Below-Basic students was necessary to harmonize proficiency categories across states but restricted the criterion range, which can reduce observed correlations (Mendoza and Mumford, 1987). The findings therefore apply only to students represented in the Basic, Proficient, and Advanced categories.

Second, the Grade 7/Grade 8 comparison is cross-sectional. Differences in sample size, score distributions, assignment mix, or cohort composition may explain the stronger Grade 8 associations; longitudinal data are needed before interpreting this pattern developmentally.

Third, some EFA cohorts were small relative to the 54 analytic variables, and one cohort had KMO=.52. Low Tucker congruence may therefore reflect estimation instability as well as genuine contextual variation (MacCallum et al., 1999). Larger samples and confirmatory or resampling-based analyses are needed before defining reusable latent scales.

Finally, gender analyses were limited to binary roster labels and do not constitute comprehensive fairness evidence; fairness in educational measurement concerns the validity and comparability of score interpretations and uses across relevant populations, rather than any single demographic comparison (American Educational Research Association et al., 2014; Johnson et al., 2022). Future work should examine other relevant student characteristics where ethically and legally appropriate.

7 Conclusion

Keystroke process data provide information about writing behavior that is not visible in final text alone, but their value depends on validation across contexts. Across two middle schools and thousands of classroom documents, Grade 8 proficiency associations involving engagement and temporal fluency replicated across schools and writing contexts, whereas Grade 7 associations were sparse, latent structures were unstable, and gender-linked patterns were context dependent. These findings support cautious use of keystroke measures in AI-enabled formative assessment and reinforce a central requirement: behavioral process signals should demonstrate validity, generalizability, and fairness before being used for automated inference or feedback.

References

American Educational Research Association, American Psychological Association, and National Council on Measurement in Education. 2014. *Standards for Educational and Psychological Testing*. American Educational Research Association, Washington, DC.

Veerle M. Baaijen, David Galbraith, and Kees de Glopper. 2012. Keystroke analysis: Reflections on procedures and measures. *Written Communication*, 29(3):246–277.

Maurice S Bartlett. 1950. Tests of significance in factor analysis. *British Journal of Statistical Psychology*, 3(2):77–85.

Yoav Benjamini and Yosef Hochberg. 1995. Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B*, 57(1):289–300.

Norman Cliff. 1993. Dominance statistics: Ordinal analyses to answer ordinal questions. *Psychological Bulletin*, 114(3):494–509.

Rianne Conijn, Christine Cook, Menno van Zaanen, and Luuk Van Waes. 2022. Early prediction of writing quality using keystroke logging. *International Journal of Artificial Intelligence in Education*, 32:835–866.

Rianne Conijn, Jens Roeser, and Menno van Zaanen. 2019. Understanding the keystroke log: The effect of writing task on keystroke features. *Reading and Writing*, 32(9):2353–2374.

Paul Deane, Neil Odendahl, Thomas Quinlan, Mary Fowles, Cindy Welsh, and Jennifer Bivens-Tatum. 2008. Cognitive models of writing: Writing proficiency as a complex integrated skill. *ETS Research Report Series*, 2008(2):i–36.

Paul Deane, Mo Zhang, Jiangang Hao, and Chen Li. 2026. Using keystroke dynamics to detect nonoriginal text. *Journal of Educational Measurement*, 63(1):e12431.

Linda Flower and John R. Hayes. 1981. A cognitive process theory of writing. *College Composition and Communication*, 32(4):365–387.

Steve Graham, Karen R. Harris, Charles MacArthur, and Tanya Santangelo. 2017. Self-regulation and writing. In *Handbook of Self-Regulation of Learning and Performance*, pages 138–152. Routledge.

Steve Graham, Michael Hebert, and Karen R. Harris. 2015. Formative assessment and writing: A meta-analysis. *The Elementary School Journal*, 115(4):523–547.

John L. Horn. 1965. A rationale and test for the number of factors in factor analysis. *Psychometrika*, 30(2):179–185.

Paul Jaccard. 1901. Étude comparative de la distribution florale dans une portion des alpes et des jura. *Bulletin de la Société Vaudoise des Sciences Naturelles*, 37:547–579.

Yang Jiang, Mo Zhang, Jiangang Hao, Paul Deane, and Chen Li. 2024. Using keystroke behavior patterns to detect nonauthentic texts in writing assessments: Evaluating the fairness of predictive models. *Journal of Educational Measurement*, 61(4):571–594.

Matthew S. Johnson, Xiang Liu, and Daniel F. McCaffrey. 2022. Psychometric methods to evaluate measurement and algorithmic bias in automated scoring. *Journal of Educational Measurement*, 59(3):338–361.

Henry F. Kaiser. 1960. The application of electronic computers to factor analysis. *Educational and Psychological Measurement*, 20(1):141–151.

Mariëlle Leijten and Luuk Van Waes. 2013. Keystroke logging in writing research: Using Inputlog to analyze and visualize writing processes. *Written Communication*, 30(3):358–392.

Diane Litman, Tazin Afrin, Omid Kashefi, Christopher Olshefski, Amanda Godley, and Rebecca Hwa. 2022. An automated writing evaluation system for supporting self-monitored revising. In *International Conference on Artificial Intelligence in Education*, pages 581–587. Springer.

Urbano Lorenzo-Seva and Jos M. F. ten Berge. 2006. Tucker's congruence coefficient as a meaningful index of factor similarity. *Methodology*, 2(2):57–64.

Robert C. MacCallum, Keith F. Widaman, Shaobo Zhang, and Sehee Hong. 1999. Sample size in factor analysis. *Psychological Methods*, 4(1):84–99.

Jorge L. Mendoza and Michael D. Mumford. 1987. Corrections for attenuation and range restriction on the predictor. *Journal of Educational Statistics*, 12(3):282–293.

Neil Odendahl and Paul Deane. 2018. Assessing the writing process: A review of current practice. Technical Report ETS-RM-18-10, Educational Testing Service.

Andrea Révész, Hyeonjeong Jeong, Yuyong Liu, Tsubasa Otsuki, Xin Rong, Ting Zeng, Makoto Takahashi, and Motoaki Sugiura. 2026. The neural correlates of pausing behaviours in l2 writing. *Journal of Second Language Writing*, 72:101303.

Valerie J. Shute, Lihua Wang, Samuel Greiff, Wei Zhao, and Gregory Moore. 2016. Measuring problem solving skills via stealth assessment in an engaging video game. *Computers in Human Behavior*, 63:106–117.

Luuk Van Waes and Mariëlle Leijten. 2015. Fluency in writing: A multidimensional perspective on writing fluency applied to L1 and L2. *Computers and Composition*, 38:79–95.

Mo Zhang, Paul Deane, Andrew Hoang, Hongwen Guo, and Chen Li. 2025. Applications and modeling of keystroke logs in writing assessments. *Educational Measurement: Issues and Practice*, 44(2):5–19.

Barry J. Zimmerman and Rafael Risemberg. 1997. Becoming a self-regulated writer: A social cognitive perspective. *Contemporary Educational Psychology*, 22(1):73–101.

A Detailed Results

A.1 RQ1: Feature–Proficiency Associations

Tables 6–9 report significant feature–proficiency associations after FDR correction. Here, G denotes grade and p_{adj} denotes the FDR-adjusted p-value.

G	Feature	ρ	p_{adj}
8	overallReadingTime	0.3481	0.0242
8	stdLogBetweenWordInsertTime	0.3366	0.0242
8	varLogClauseBreakInsertTime	0.3024	0.0474
8	overallMenuTime	0.2927	0.0474
8	nBetweenWordInserts	0.2887	0.0474
8	stdLogMinorPunctInsertTime	0.2823	0.0480
8	stdLogWordStartInsertTime	0.2709	0.0575
8	overallActiveWritingTime	0.2504	0.0770
8	varLogBetweenWordInsertTime	0.2497	0.0770
8	meanLogClauseBreakInsertTime	-0.2440	0.0770
8	overallNavigationTime	0.2437	0.0770
8	medianLogClauseBreakInsertTime	-0.2344	0.0893

Table 6: Significant feature–ELA proficiency correlations for School 1 in the ELA writing context.

G	Feature	ρ	p_{adj}
7	overallActiveWritingTime	0.2931	0.0546
7	overallMenuTime	0.2838	0.0546
7	nBetweenWordInserts	0.2561	0.0947
8	nBetweenWordInserts	0.3851	0.0024
8	overallReadingTime	0.3811	0.0024
8	overallActiveWritingTime	0.3327	0.0119
8	stdLogBetweenWordInsertTime	0.2953	0.0286
8	overallMenuTime	0.2938	0.0286
8	varLogWordStartInsertTime	0.2829	0.0341
8	overallNavigationTime	0.2622	0.0555
8	varLogBetweenWordInsertTime	0.2438	0.0764
8	varLogInWordInsertTime	0.2423	0.0764
8	varLogStartingBackspaceTime	0.2337	0.0871
8	varLogClauseBreakInsertTime	0.2283	0.0913

Table 7: Significant feature–ELA proficiency correlations for School 1 in the Full writing dataset.

G	Feature	ρ	p_{adj}
8	medianLogClauseBreakInsertTime	-0.4328	0.0000
8	overallNavigationTime	0.4067	0.0000
8	meanLogClauseBreakInsertTime	-0.4007	0.0000
8	overallReadingTime	0.3794	0.0000
8	overallActiveWritingTime	0.3515	0.0000
8	nBetweenWordInserts	0.3334	0.0002
8	meanLogInWordInsertTime	-0.2838	0.0027
8	medianLogWordStartInsertTime	-0.2709	0.0043
8	medianLogInWordInsertTime	-0.2629	0.0055
8	varLogStartingBackspaceTime	0.2518	0.0080
8	varLogContinuingBackspaceTime	0.2417	0.0111
8	meanLogWordStartInsertTime	-0.2282	0.0171
8	overallMenuTime	0.2267	0.0171
8	medianLogEditTime	-0.2187	0.0216
8	varLogEditTime	0.2123	0.0255
8	meanLogContinuingBackspaceTime	-0.2096	0.0264
8	medianLogContinuingBackspaceTime	0.2004	0.0344
8	medianLogStartingBackspaceTime	-0.1983	0.0350
8	meanLogEditTime	-0.1832	0.0546
8	meanLogStartingBackspaceTime	-0.1814	0.0551
8	varLogClauseBreakInsertTime	0.1774	0.0594
8	stdLogMinorPunctInsertTime	-0.1703	0.0705

Table 8: Significant feature–ELA proficiency correlations for School 2 in the ELA writing context.

G	Feature	ρ	p_{adj}
7	overallNavigationTime	0.3974	0.0622
8	overallReadingTime	0.4735	0.0000
8	nBetweenWordInserts	0.4093	0.0000
8	overallNavigationTime	0.3961	0.0000
8	overallActiveWritingTime	0.3858	0.0000
8	medianLogClauseBreakInsertTime	-0.3825	0.0000
8	medianLogWordStartInsertTime	-0.3414	0.0001
8	meanLogClauseBreakInsertTime	-0.3399	0.0001
8	meanLogWordStartInsertTime	-0.3026	0.0006
8	meanLogInWordInsertTime	-0.3016	0.0007
8	medianLogInWordInsertTime	-0.2975	0.0007
8	medianLogStartingBackspaceTime	-0.2660	0.0032
8	stdLogBeforeWordInsertTime	-0.2629	0.0034
8	varLogContinuingBackspaceTime	0.2608	0.0034
8	meanLogStartingBackspaceTime	-0.2507	0.0050
8	overallMenuTime	0.2407	0.0072
8	stdLogInWordInsertTime	-0.1920	0.0437
8	stdLogMinorPunctInsertTime	-0.1840	0.0525
8	stdLogEditTime	-0.1831	0.0525
8	medianLogBetweenWordInsertTime	-0.1619	0.0969

Table 9: Significant feature–ELA proficiency correlations for School 2 in the Full writing dataset.

A.2 RQ2: Illustrative Factor Solutions

The following tables present the Grade 8 ELA factor solutions for both schools as illustrative examples of the recurring behavioral dimensions and variation in their composition across schools.

Factor	Dominant features	Theme
F1	overallNavigationTime, overallActiveWritingTime, overallReadingTime, overallMenuTime, varLogMinorPunctInsertTime, varLogBetweenWordInsertTime, varLogStartingBackspaceTime, varLogInWordInsertTime	Global engagement time
F2	medianLogInWordInsertTime, meanLogStartingBackspaceTime, medianLogStartingBackspaceTime, medianLogBetweenWordInsertTime, meanLogBetweenWordInsertTime, meanLogWordStartInsertTime, medianLogWordStartInsertTime	Quick revision
F3	meanLogContinuing BackspaceTime, meanLogEditTime, medianLogEditTime, medianLogContinuingBackspaceTime, stdLogContinuingBackspaceTime	Sustained revision

Table 10: Dominant features and interpreted factor themes for School 1, Grade 8, ELA writing.

Table 10 continued

Factor	Dominant features	Theme
F4	stdLogClauseBreakInsertTime, stdLogBetweenWordInsertTime, stdLogInWordInsertTime, stdLogStartingBackspaceTime, stdLogWordStartInsertTime	Edit variability
F5	medianLogMinorPunctInsertTime	Minor-punctuation rhythm
F6	meanLogBeforeWordInsertTime, medianLogBeforeWordInsertTime	Pre-word planning

Factor	Dominant features	Theme
F1	meanLogInWordInsertTime, medianLogInWordInsertTime, meanLogBetweenWordInsertTime, medianLogBetweenWordInsertTime, meanLogStartingBackspaceTime, medianLogStartingBackspaceTime, meanLogWordStartInsertTime, medianLogWordStartInsertTime, meanLogClauseBreakInsertTime, medianLogClauseBreakInsertTime	Quick revision / clause-level rhythm
F2	overallActiveWritingTime, overallNavigationTime, overallReadingTime, overallMenuTime, nBetweenWordInserts, varLogEditTime, varLogStartingBackspaceTime, varLogContinuingBackspaceTime	Global engagement time
F3	meanLogContinuingBackspaceTime, meanLogEditTime, medianLogEditTime, medianLogContinuingBackspaceTime	Sustained revision
F4	stdLogInWordInsertTime, stdLogBetweenWordInsertTime, stdLogWordStartInsertTime, stdLogBeforeWordInsertTime, stdLogMinorPunctInsertTime, stdLogStartingBackspaceTime	Edit variability
F5	meanLogBeforeWordInsertTime, stdLogClauseBreakInsertTime	Pre-word planning
F6	medianLogMinorPunctInsertTime, meanLogMinorPunctInsertTime	Minor-punctuation rhythm

Table 11: Dominant features and interpreted factor themes for School 2, Grade 8, ELA writing.

A.3 RQ3: Gender-Linked Feature Differences

The following tables report feature-level gender differences that remained significant after FDR correction. Negative Cliff's δ denotes higher values for female students, whereas positive values denote higher values for male students. No significant differences were observed for School 1 Grade 8 or for School 2 Grade 7 in the ELA writing context.

Feature	p_{adj}	Cliff's δ
nBetweenWordInserts	$< .0001$	-0.5710
stdLogClauseBreakInsertTime	0.0003	0.5249
overallNavigationTime	0.0003	-0.5164
overallActiveWritingTime	0.0003	-0.5062
overallReadingTime	0.0005	-0.4832
varLogContinuingBackspaceTime	0.0008	-0.4670
stdLogBetweenWordInsertTime	0.0033	0.4132
overallMenuTime	0.0033	-0.4128
varLogInWordInsertTime	0.0098	-0.3710
varLogEditTime	0.0166	-0.3390
varLogBetweenWordInsertTime	0.0186	-0.3394
stdLogBeforeWordInsertTime	0.0295	0.3168
meanLogMinorPunctInsertTime	0.0295	-0.3117
varLogWordStartInsertTime	0.0295	-0.3096
varLogStartingBackspaceTime	0.0443	-0.2913

Table 12: Significant gender-linked feature differences for School 1, Grade 7, ELA writing (35 male, 67 female).

Feature	p_{adj}	Cliff's δ
nBetweenWordInserts	$< .0001$	-0.5657
stdLogBetweenWordInsertTime	$< .0001$	0.5429
overallNavigationTime	0.0009	-0.4585
stdLogClauseBreakInsertTime	0.0009	0.4514
overallActiveWritingTime	0.0020	-0.4214
stdLogBeforeWordInsertTime	0.0034	0.4007
overallMenuTime	0.0095	-0.3629
stdLogInWordInsertTime	0.0142	0.3443
overallReadingTime	0.0185	-0.3307
meanLogMinorPunctInsertTime	0.0382	-0.2993
medianLogWordStartInsertTime	0.0503	0.2836
meanLogWordStartInsertTime	0.0503	0.2807
stdLogContinuingBackspaceTime	0.0706	0.2629
varLogStartingBackspaceTime	0.0964	-0.2446
stdLogEditTime	0.0975	0.2407
medianLogMinorPunctInsertTime	0.0975	-0.2332
varLogWordStartInsertTime	0.0975	-0.2243
varLogEditTime	0.0975	-0.2021

Table 13: Significant gender-linked feature differences for School 1, Grade 7, Full writing dataset (40 male, 70 female).

Feature	p_{adj}	Cliff's δ
medianLogWordStartInsertTime	0.0264	0.5097
meanLogWordStartInsertTime	0.0593	0.4452

Table 14: Significant gender-linked feature differences for School 2, Grade 7, Full writing dataset (30 male, 31 female).

Feature	p_{adj}	Cliff's δ
meanLogBeforeWordInsertTime	0.0591	0.2782
nBetweenWordInserts	0.0591	-0.2742
medianLogBeforeWordInsertTime	0.0591	0.2699
overallNavigationTime	0.0591	-0.2562
overallReadingTime	0.0591	-0.2523
overallActiveWritingTime	0.0591	-0.2494
varLogEditTime	0.0882	-0.2272
medianLogClauseBreakInsertTime	0.0882	0.2263

Table 15: Significant gender-linked feature differences for School 2, Grade 8, ELA writing (75 male, 74 female).

Feature	p_{adj}	Cliff's δ
stdLogClauseBreakInsertTime	0.0117	0.3144
nBetweenWordInserts	0.0117	-0.3139
overallNavigationTime	0.0117	-0.3114
stdLogBeforeWordInsertTime	0.0240	0.2844
overallMenuTime	0.0348	-0.2655
overallActiveWritingTime	0.0348	-0.2618
overallReadingTime	0.0521	-0.2432
stdLogBetweenWordInsertTime	0.0521	0.2402
medianLogClauseBreakInsertTime	0.0827	0.2208
medianLogBeforeWordInsertTime	0.0992	0.2106

Table 16: Significant gender-linked feature differences for School 2, Grade 8, Full writing dataset (79 male, 75 female).

Self-Revising Agents as Item Writers: Benchmarking Against Professional Item Writing

Steven Tang
eMetric LLC
steven@emetric.net

Zhen Li
eMetric LLC
zli@emetric.net

Richard Patz
University of California, Berkeley
rpatz@berkeley.edu

Abstract

Expert reviewers rated 156 calculus items from a self-revising multi-agent framework (Claude Sonnet 4.6 or GPT-5.4, three prompt conditions) and 26 human-written items. AI items were formative-ready nearly as often as human items (78–80% versus 88%) but summative-ready less often (26% and 10% versus 58%); a root-item reference mattered most.

1 Background

In the recent past, Automatic Item Generation (AIG) has relied on rule-based templates (Gierl and Haladyna, 2013). Contemporary Large Language Models (LLM) can now produce contextually rich items from natural-language prompts alone (Attali et al., 2022; Li et al., 2025). Li et al. (2025) found that approximately 25% of GPT-4-generated clinical items were free of major flaws and roughly half were useful starting points relative to writing from scratch.

However, when generating items from a standard alone, the model must infer content, format, difficulty, and style without guidance, and outputs often drift from the conventions of a specific course or program. Grounding generation in existing professionally written items addresses this limitation: supplying them as in-context exemplars (Brown et al., 2020) or through retrieval-grounded prompting (Lewis et al., 2020) steers the model toward established conventions and transfers design decisions that are difficult to articulate fully in a prompt.

This study investigates whether a self-revising multi-agent (SRMA) AI framework can produce useful calculus items under different prompt conditions, including conditions where a root item or a root item paired with a derived variant is provided as reference. The root items used in this study were produced by professional item writers during the 2020 pandemic whereby existing hand-written exam materials were carefully converted into standard-aligned digital items. To support the potential of multiple forms, the item writers then created "derived" variant items from the root items. This process of creating derived items from root materials is now a seemingly natural candidate for AI to address.

2 Aims

Aim 1: Evaluate whether an SRMA pipeline can expand a bank of calculus items from human-constructed base materials at a quality comparable to professional item writers.

Aim 2: Determine how much reference material the pipeline needs, by comparing items generated from the content standard alone, from a root item, and from a root item paired with one human-written variant, and whether the second exemplar adds anything beyond the root.

Aim 3: Examine how Claude and GPT, used as backends in the pipeline, differ in output verbosity, mathematical precision, and adherence to the base item.

3 Samples

The dataset consisted of 78 items authored by professional item writers for a university-level calculus course, mapped across 6 content strands such as Functions, Limits, Derivatives, and Integrals. Items were organized into 26 item families, each anchored by a root item and two human-authored derived variants. The item writers were asked to expand the bank from the root items to support multiple test forms by creating clones or variants of each root; the guidance did not specify which. Sixteen families were multiple-choice, five were fill-in-the-blank, and five were open-ended.

4 Methods

4.1 Self-Revising Multi-Agent Pipeline

Items were generated by a self-revising multi-agent (SRMA) pipeline (Figure 1). Each item began at a

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 369–376
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

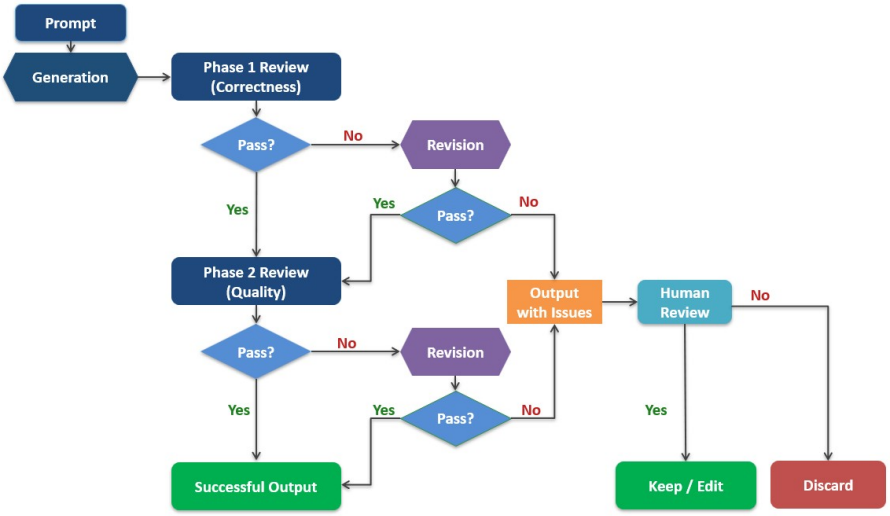


Figure 1: Diagram of SRMA Item Generation Pipeline.

Generation stage, where an LLM received a structured prompt encoding the target standard, item-writing guidelines, and reference examples from existing human-authored items. The model returned a draft item in structured format, which then entered two successive Review phases.

Phase 1 (Correctness) evaluated whether the item was mathematically sound and properly aligned to its target standard. Issues that could be resolved automatically were fixed programmatically; others triggered a targeted LLM revision followed by re-evaluation. Phase 2 (Quality) evaluated the passing items on dimensions including distractor plausibility, structural compliance, and language clarity, with the same revision-and-re-evaluate loop applied to any failures. Generated items were serialized to QTI/XML. Items failing revision in each phase were flagged for human review with a summary of remaining issues. Two generated items had a failed figure generation; these two items were discarded and replaced with a second generation attempt, which succeeded in producing the figure.

4.2 Research Design

The study employed a 2 (LLM: Claude Sonnet 4.6, GPT-5.4) × 3 (prompt condition: P1, P2, P3) fully crossed design, with human-generated items serving as a reference baseline. Two LLMs were evaluated: GPT-5.4 (OpenAI) and Claude Sonnet 4.6 (Anthropic), both demonstrating strong mathematical proficiency on established benchmarks (Wei et al., 2025; Anthropic, 2026). Three prompt conditions varied the amount of reference information provided to each LLM, reflecting increasing levels of contextual guidance:

- P1 (Standard-Only): The LLM receives the content standard and generation instructions only, with no reference item. This condition tests the model's ability to generate a standard-aligned calculus item from scratch, without any example of what a calculus item should look like.
- P2 (Root Item): Given a root item, AI generates derived items using each LLM, enabling direct comparison between AI-generated and human-generated derived items anchored to the same root.
- P3 (Item Family): The LLM receives the content standard, a root item, and a human-generated derived item from the same family. This condition provides the fullest context, showing the model what good variation within a family looks like.

As shown in Table 1, each LLM was applied under all three prompt conditions, yielding 6 AI-generation cells of 26 items each (156 AI-generated items total). Of the 52 human-generated derived variants included, 26 (one per root item) served as the comparison baseline across conditions, and the remaining 26 (one per root item) were used as family exemplars in the P3 prompt. Generated items used the same format as their root item.

The number of items per condition was determined by the number of root items available in the human-authored calculus pool. 26 items were designated as root items, each having at least two human-authored derived items.

Prompt Condition	Source		
	Claude Sonnet 4.6	GPT-5.4	Human Expert Generated
P1	26	26	52 *(26 used as comparison, 26 used as P3 exemplars)*
P2	26	26	
P3	26	26	

Table 1: Conditions for item generation.

4.3 Item Evaluation

The 156 AI-generated items and the 26 human-authored comparison items (182 items in total) were rated on the 10 criteria in Appendix A: standard alignment, course appropriateness, mathematical correctness, distractor quality (multiple-choice items only), figure quality (items with a figure only), stem clarity, stem conciseness, relationship to the root item (unrelated, variant/alternate, or clone), and two holistic judgments of how much revision the item would need before use on a summative test and on a formative test (major revision, minor revision, or no revision). Reviewers also recorded open-ended comments on each item.

Every item was rated independently by two expert reviewers: the first author, and a professional item writer experienced in item review and summative test development.[1] Agreement was summarized by exact agreement, Cohen's κ (Cohen, 1960), and Gwet's AC1 (Gwet, 2008); the three measures respond differently when ratings concentrate in one category, as they do for several criteria here, so all three are reported. Every criterion on which the reviewers differed was then resolved by the first author in consultation with the second reviewer in an adjudication review. The final score on each criterion is the agreed rating or, where the reviewers differed, the adjudicated rating. A frontier LLM (Claude Fable 5.1) also rated all 182 items with the same rubric; its ratings were not used in the final scores and are reported only in the agreement analysis.

Five item-level indicators were derived from the final scores: a *fatal flaw* (a score of 0 on standard alignment, course appropriateness, or mathematical correctness); *acceptable on all criteria* (a score of 1 on alignment, appropriateness, correctness, clarity, and conciseness, and no score of 0 on distractor or figure quality where applicable; the two holistic revision judgments are not part of this indicator); *summative-ready* and *formative-ready* (no revision needed for the respective use); and *clone* (interchangeable with the root item). A *composite quality* index averages the seven quality criteria, each rescaled to 0–1, excluding relationship to root and the two revision judgments, which are outcomes rather than components.

Because every family contributed one item to each condition, comparisons were paired by family: Wilcoxon signed-rank tests (Wilcoxon, 1945) on family means for pairwise comparisons between sources or conditions, Friedman tests (Friedman, 1937) across the three prompt conditions with families as blocks, and Fisher's exact tests (Fisher, 1922) for item-level rates. With 26 families the study can detect only large effects; p-values are reported descriptively without correction for multiple comparisons.

5 Results

5.1 Reviewer Agreement and Adjudication

Table 2 summarizes agreement between the two reviewers on each criterion. On the objective criteria (standard alignment, course appropriateness, and mathematical correctness) the reviewers agreed on 96–99% of items (AC1 = .95–.99). Agreement was lower on the judgment-laden criteria: 88% on conciseness, 80% on relationship to root (κ = .62) and on the formative verdict, 74% on distractor quality, and only 49% on the summative verdict (κ = .15, AC1 = .30). The summative verdict alone accounted for 93 of the 248 criterion-level disagreements, likely reflecting the two reviewers' different levels of experience in the professional authoring of summative assessments. The three agreement measures diverge where ratings concentrate in one category. On course appropriateness the reviewers disagreed on only one of the 182 items, giving 99.5% exact agreement; but because both rated nearly every item appropriate, κ, which discounts the agreement expected under that concentration, is .00, whereas AC1 applies a smaller correction and is .99. Across criteria AC1 tracks percent agreement closely without matching it. Composite quality indices computed from each reviewer's own ratings were identical for 62% of items and within .20 of each other for 96%; on average the two indices were .06 apart on the 0–1 scale, with reviewer means of .96 and .94.

[1] Janet Burnham, JLB Consulting, burnhams611@gmail.com.

Criterion	n	Agree (%)	κ	AC1	Splits
Standard alignment	182	95.6	.31	.95	8
Course appropriateness	182	99.5	.00	.99	1
Mathematical correctness	182	96.7	.39	.97	6
Distractor quality	111	73.9	.12	.70	29
Figure quality	20	80.0	−.08	.78	4
Stem clarity	182	92.9	−.03	.92	13
Stem conciseness	182	87.9	.37	.85	22
Relationship to root	182	80.2	.62	.73	36
Revision – summative	182	48.9	.15	.30	93
Revision – formative	182	80.2	.16	.78	36

Table 2: Agreement between the two expert reviewers by criterion (splits = disagreements resolved in adjudication).

Adjudication was conservative. On the evaluative criteria, 179 of the 212 disagreements (84%) were resolved to the more critical of the two ratings, including 84 of the 93 summative-verdict splits, 30 of 36 formative splits, 26 of 29 distractor splits, and all 22 conciseness splits. Relationship to root had 31 of 36 splits resolved toward being more similar to the root.

Compared with the final adjudicated scores, the AI reviewer matched 85% of all criterion-level judgments: 98–99% on standard alignment, course appropriateness, and mathematical correctness; 93% on stem clarity; 88% on conciseness; 82% on the formative verdict; 76% on relationship to root; 70% on distractor quality; and 56% on the summative verdict. Its profile thus mirrored the human reviewers': near-perfect on the objective criteria and weakest on the summative verdict.

5.2 Item Quality by Source

Table 3 reports the final scores by source. Human-authored variants had the highest composite quality (.98), followed by Claude (.94; p = .064) and GPT (.92; p = .006). Fatal flaws were rare for every source: 1 of 26 human items (3.8%), 3 of 78 Claude items (3.8%), and 2 of 78 GPT items (2.6%). The flaws differed in kind, however. The one human item with a fatal flaw was misaligned, in that it did not require the process named in the standard, whereas four of the five AI items with a fatal flaw had a mathematical flaw (a missing or duplicated key, or an ill-posed function) and one was misaligned.

	Human (n=26)	Claude (n=78)	GPT (n=78)
Criterion mean			
Standard alignment (0–1)	.96	1.00	.99
Course appropriateness (0–1)	1.00	1.00	1.00
Mathematical correctness (0–1)	1.00	.96	.99
Distractor quality (0–2)[a]	1.80	1.71	1.69
Figure quality (0–2)[b]	2.00	1.62	1.89
Stem clarity (0–1)	.96	.96	.91
Stem conciseness (0–1)	1.00	.82	.79
Relationship to root (0–2)	1.62	1.64	1.41
Revision – summative (0–2)	1.54	1.18	1.09
Revision – formative (0–2)	1.88	1.72	1.78
Composite quality (0–1)	.98	.94	.92
Indicator (% of items)			
Fatal flaw	3.8	3.8	2.6
Acceptable on all criteria	92.3	75.6	70.5
Summative-ready	57.7	25.6	10.3
Formative-ready	88.5	78.2	79.5
Clone of root	69.2	64.1	42.3

Table 3: Final scores by source. [a] Multiple-choice items only (n = 15, 48, 48). [b] Items with a figure only (n = 3, 8, 9).

The share of items acceptable on all criteria was 92% for human items, 76% for Claude (Fisher's exact p = .09), and 71% for GPT (p = .03). The criterion means in Table 3 locate the differences: AI items were less concise (.82 and .79 versus 1.00), had slightly weaker distractor sets (1.71 and 1.69 versus 1.80), and GPT stems were less often clear (.91 versus .96), whereas alignment was near ceiling for every source (.96–1.00). Conciseness accounted for most of the shortfall: of the 42 AI items that did not meet every criterion, 26 missed on conciseness alone and 4 more on conciseness together with an unclear stem; the remaining 12 involved an unclear stem (5), a mathematical flaw (4), defective distractors (2), or misalignment (1). Claude's items were as close to their roots as the human variants (64% versus 69% clones), while GPT's were farther (42%). These comparisons pool the three prompt conditions; Section 5.3 compares the root-conditioned prompts alone with the human variants.

The sources diverged sharply on readiness. 58% of human variants were judged usable on a summative test without revision, against 26% of Claude items (p = .008) and 10% of GPT items (p < .001). Major revision was required for 6 Claude items, 1 GPT item, and 1 human item. For formative use the gap largely disappeared: 78% (Claude), 80% (GPT), and 88% (human) of items

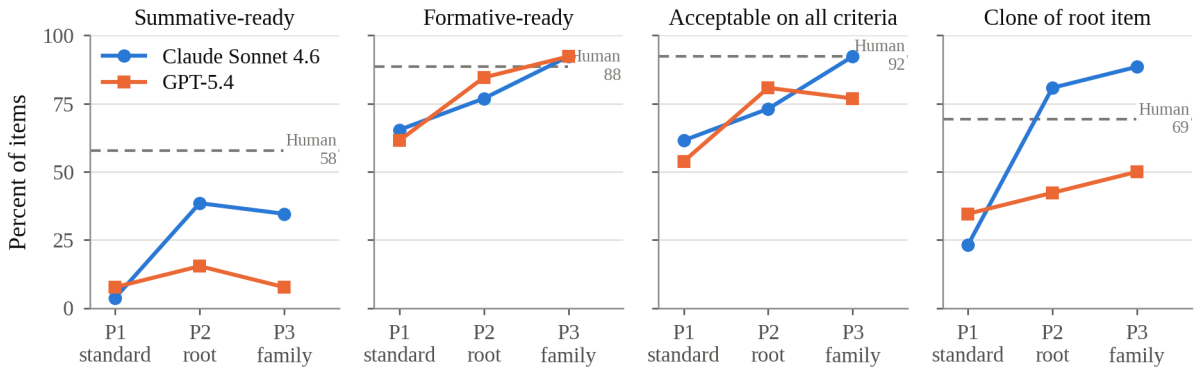


Figure 2: Summative readiness, formative readiness, share of items acceptable on all criteria, and clone rate by prompt condition and model (n = 26 items per point; dashed lines mark the human-written variants).

needed no revision, and none of the pairwise differences was statistically significant (all p ≥ .24).

5.3 Effect of Reference Material

Of the two design factors, the amount of reference material had the larger effect, moving every indicator more than the choice of model did (Table 4, Figure 2). Pooled over the two models, moving from the standard alone (P1) to a root item (P2) to an item family (P3) raised the share of items acceptable on all criteria from 58% to 77% to 85%, formative readiness from 63% to 81% to 92% (Friedman p < .001), the clone rate from 29% to 62% to 69% (p < .001), and mean conciseness from .67 to .85 to .90 (p = .003).

Condition	Composite	Fatal	All accept.	Summ. ready	Form. ready	Clone	Concise
Claude P1	.91	0.0	61.5	3.8	65.4	23.1	.65
Claude P2	.94	7.7	73.1	38.5	76.9	80.8	.85
Claude P3	.97	3.8	92.3	34.6	92.3	88.5	.96
GPT P1	.89	3.8	53.8	7.7	61.5	34.6	.69
GPT P2	.94	0.0	80.8	15.4	84.6	42.3	.85
GPT P3	.94	3.8	76.9	7.7	92.3	50.0	.85
Human	.98	3.8	92.3	57.7	88.5	69.2	1.00

Table 4: Indicators by generation condition (n = 26 items per row; All accept. = acceptable on all criteria).

For Claude, a root item transformed the output: summative readiness rose from 4% in P1 to 38% in P2 and 35% in P3 (p = .002), the clone rate from 23% to 81% to 88% (p < .001), the share acceptable on all criteria from 62% to 73% to 92%, and formative readiness from 65% to 77% to 92% (p = .025). For GPT, reference material improved the share acceptable on all criteria (54%, 81%, 77%), composite quality (.89, .94, .94; p = .031), and formative readiness (62%, 85%, 92%; p = .009) but not summative readiness (8%, 15%, 8%; p = .37), and its items remained farther from the root (35%, 42%, 50% clones; p = .22).

Two further patterns matter for practice. First, the second exemplar improved criterion-level quality rather than readiness: from P2 to P3, the share of Claude's items acceptable on all criteria rose from 73% to 92% and its formative readiness from 77% to 92%, but its summative readiness did not (38% versus 35%). Second, reference material did nothing for mathematical correctness: fatal flaws occurred at a low rate in every condition (pooled 1.9%, 3.8%, 3.8%; Claude 0%, 7.7%, 3.8%; GPT 3.8%, 0%, 3.8%), and all three of Claude's mathematical errors occurred in items generated with a reference item. In the strongest AI condition, Claude with an item family, 92% of items were acceptable on all criteria, 88% were clones of their root, and 92% were ready for formative use, yet only 35% were ready for summative use, compared with 58% of the human variants.

The human variants also provide a like-for-like comparison for the root-conditioned prompts. The item writers and the P2 and P3 runs started from the same root items; the writers were asked for clones or variants of each root, the prompts asked for an item variant, and neither specified how close to the root the new item should stay. The comparison therefore shows where professionals and models each landed on the clone–variant spectrum given the same starting material. With the root supplied, Claude's items were classified as clones at least as often as the human variants (81% in P2 and 88% in P3 versus 69%; p = .37 and .10), while GPT's items departed from the root more often (42% and 50% clones; p = .07 and .13). On the readiness verdicts, Claude's items were summative-ready less often

than the human variants (38% and 35% versus 58%), a difference that was not statistically significant (p = .13 and .11), whereas GPT's items were summative-ready much less often (15% and 8%; p = .002 and < .001); formative readiness was at the human level for both models in P3 (92% versus 88%).

5.4 Model Differences

Compared directly with each other (Aim 3), the two models did not differ in composite quality, .94 for Claude and .92 for GPT (p = .38), but their profiles differed in several respects. Claude's items were summative-ready more often than GPT's, 26% against 10% over all conditions (p = .017), and the difference held within the two root-conditioned prompts, 38% against 15% in P2 (p = .034) and 35% against 8% in P3 (p = .020). Claude's items were also closer to the root: 64% were clones, against 42% of GPT's (p = .011). GPT's items, in turn, contained fewer mathematical errors, 1 against Claude's 3, and fewer required major revision, 1 against 6 (p = .12). At the same time, GPT's stems were more often unclear, 7 against 3, and drew more verbosity comments from the reviewers, 42 against 27 mentions per 100 items summed over the three raters. On formative readiness the models were indistinguishable, at 78% and 80%.

5.5 Item Format

The differences reported above depended on item format (Table 5). For multiple-choice items, both models approached the human benchmark on formative readiness (Claude 85%, GPT 94%, human 93%) and on the share acceptable on all criteria (85%, 90%, 100%). For open-ended items, AI performance fell away: formative readiness dropped to 53% (Claude) and 47% (GPT) against 83% for human items, the share acceptable on all criteria to 47% and 33%, and mean conciseness to .47 and .40; two of Claude's major-revision items were open-ended items generated from the standard alone. Fill-in-the-blank items separated the models, with Claude reaching 80% formative and 47% summative readiness against GPT's 67% and 13%. These cells are small (15 AI items and 5–6 human items per format) and were not tested.

Format / source	n	All accept.	Summ.-ready	Form.-ready
Multiple choice				
Human	15	100.0	60.0	93.3
Claude	48	85.4	22.9	85.4
GPT	48	89.6	12.5	93.8
Fill-in-the-blank				
Human	5	80.0	80.0	80.0
Claude	15	73.3	46.7	80.0
GPT	15	46.7	13.3	66.7
Open-ended				
Human	6	83.3	33.3	83.3
Claude	15	46.7	13.3	53.3
GPT	15	33.3	0.0	46.7

Table 5: Item-level indicators by item format and source (percent of items; All accept. = acceptable on all criteria).

6 Discussion

With respect to Aim 1, the results indicate that the pipeline can contribute usable items to a calculus bank. AI items were as free of fatal flaws as the professional variants (3–4% for all sources), were acceptable on all criteria about three-quarters of the time, and were judged formative-ready at rates close to those of the human items (78–80% versus 88%). The gap in summative readiness consisted mainly of minor revisions, such as trimming stems and strengthening distractor sets, rather than major ones. In the conditions most comparable to the item writers' own task, the root-conditioned prompts, Claude's items were summative-ready 35–38% of the time against 58% for the human variants, a difference that was not statistically significant, whereas GPT's items remained well below (8–15%). The few mathematical errors that occurred had passed the pipeline's own review phases, an indication that AI review of AI-generated content, while useful, is not a complete substitute for human checking.

Among the factors examined, the amount of reference material had the largest effect (Aim 2). Items generated from the standard alone tended to restate definitions and scaffold the solution method, whereas items generated from a root item followed the conventions of the bank more closely. Adding a second exemplar was associated with a higher share of items acceptable on all criteria but not with higher summative readiness. The root items were written by professionals and served as good examples, but they had not been screened to be flawless, and the human variants derived from

them were judged summative-ready 58% of the time. Because generation conditions on the root, a root that would itself need revision may carry its shortcomings into the generated item, so the summative-readiness rates observed here may be lower than those attainable with roots rigorously curated for high-stakes summative purposes.

The two models differed in profile more than in overall quality (Aim 3). Claude's items were more often close to the root and more often summative-ready; GPT's items contained fewer mathematical errors but were more often verbose and less closely tied to the root. Which profile is preferable depends on whether the goal is parallel forms or broader coverage. A pattern common to both models was the inclusion of material that a summative item would omit, such as hints, restated definitions, and decorative prose. This suggests that the models' default notion of an item is closer to a formative than a summative one. Whether that default is a limitation depends on the intended use, and how prompting strategies can best be tuned to different uses remains open for future work.

References

Anthropic. 2026. Claude Sonnet 4.6 system card.

Yigal Attali, Andrew Runge, Geoffrey T. LaFlair, Kevin Yancey, Sarah Goodwin, Yena Park, and Alina A. von Davier. 2022. The interactive reading task: Transformer-based automatic item generation. *Frontiers in Artificial Intelligence*, 5:903077.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D. Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language models are few-shot learners. In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.

Jacob Cohen. 1960. A coefficient of agreement for nominal scales. *Educational and Psychological Measurement*, 20(1):37–46.

Ronald A. Fisher. 1922. On the interpretation of χ^2 from contingency tables, and the calculation of P. *Journal of the Royal Statistical Society*, 85(1):87–94.

Milton Friedman. 1937. The use of ranks to avoid the assumption of normality implicit in the analysis of variance. *Journal of the American Statistical Association*, 32(200):675–701.

Mark J. Gierl and Thomas M. Haladyna, editors. 2013. *Automatic Item Generation: Theory and Practice*. Routledge, New York, NY.

Kilem L. Gwet. 2008. Computing inter-rater reliability and its variance in the presence of high agreement. *British Journal of Mathematical and Statistical Psychology*, 61(1):29–48.

Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. Retrieval-augmented generation for knowledge-intensive NLP tasks. In *Advances in Neural Information Processing Systems*, volume 33, pages 9459–9474. Curran Associates, Inc.

Jiaxuan Li, Saed Rezayi, Peter Baldwin, Polina Harik, and Victoria Yaneva. 2025. Towards reliable generation of clinical chart items: A counterfactual reasoning approach with large language models. In *Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con): Full Papers*, pages 142–153, Pittsburgh, Pennsylvania, USA. National Council on Measurement in Education (NCME).

Hu Wei, Ze Xu, Boyu Yang, Linlin Miao, Weiqi Zhai, Yihan Li, Zixuan Li, Zhijun Wang, Boya Wang, Jianwei Yu, Jialing Yuan, Xiaoyue Zhang, Cheng He, Minglei Chen, Zifan Zhang, Qianhui Li, Wei Wang, and Xiang Xu. 2025. SKYLENAGE technical report: Mathematical reasoning and contest-innovation benchmarks for multi-level math evaluation. *Preprint*, arXiv:2510.01241.

Frank Wilcoxon. 1945. Individual comparisons by ranking methods. *Biometrics Bulletin*, 1(6):80–83.

A Item Evaluation Criteria

Table A.1 lists the criteria applied by the expert reviewers to every AI-generated and human-authored comparison item (see Section 4.3).

Criterion	What Reviewers Evaluate	Scale
Standard Alignment	Item requires the cognitive process specified in the standard or learning objective	0 = Not aligned 1 = Aligned
Course Appropriateness	Content and demand suitable for university calculus students	0 = Inappropriate 1 = Appropriate
Mathematical Correctness	The mathematical content, terms, definitions, and symbols used in the item must be entirely correct. For Multiple Choice Item: – Key is unambiguously correct; – All distractors are definitively incorrect	0 = Error present 1 = Correct
Distractor Quality (Set) (Multiple Choice items only)	Acceptable distractors are unambiguously incorrect. Good distractors are acceptable but also plausible under an identifiable misconception The set of distractors is: **Defective** if any element is not Acceptable **Acceptable** if each element is only Acceptable **Good** if each element is Acceptable, at least one is Good, and no option (including the correct one) stands out suggestively (i.e., is an 'outlier') **Excellent** if each element is Good and there are no outliers	0 = Defective 1 = Acceptable 2 = Good or Excellent
Figure Quality	Figure (if any) is mathematically correct, meaningful and displayed well.	0 = Defective 1 = Usable with minor revision 2 = Good
Stem Clarity	Stem language is clear for the math task – Question is unambiguous	0 = Unclear 1 = Appropriate
Stem Conciseness	Stem language is concise for the math task – No decorative or redundant prose	0 = Verbose 1 = Appropriate
Relationship to "Root" Item	**Clone** that is interchangeable with root (same standard, same concept, same skills, same difficulty level); **Variant/Alternate** that addresses the same standard but uses a similar or distinctly different conceptual approach, required skills, or difficulty; **Unrelated**	0 = Unrelated 1 = Variant/Alternate 2 = Clone
Open Ended	Make a few observations about this item, anything that stands out as particularly strong or in need of revision.	
Revision for Summative Use	How much revision is needed before this item could be used on a summative test?	0 = Needs major revision 1 = Needs minor revision 2 = No revision needed
Revision for Formative Use	How much revision is needed before this item could be used on a formative test?	0 = Needs major revision 1 = Needs minor revision 2 = No revision needed

Table A.1: Item evaluation criteria applied by expert human reviewers.

Evaluating Neural Network-Based IRT Estimation in Multistage Adaptive Testing with Small Item Banks

Seong Eun Hong
Center for Applied Linguistics
shong@cal.org

Abstract

Neural network (NN)-based item response theory (IRT) estimation offers a potential alternative for calibrating items when operational samples are small, but its accuracy may depend on how well simulated training data represents later applications. This study investigates the performance of NN-based IRT estimation in multistage adaptive language assessment with small item banks and short tests. Results showed that accuracy improved with larger samples and more training information. Moderate distribution shifts were largely accommodated under higher iteration conditions, whereas larger shifts continued to reduce estimation accuracy.

1 Introduction

Accurate item parameter estimation is essential for item response theory (IRT), but calibration becomes challenging when only small samples are available. Limited response information can result in unstable parameter estimates, and this becomes more challenging in adaptive testing, where responses are sparse because examinees encounter different subsets of items (Ban et al., 2002). In multistage testing (MST), routing further structures item exposure by assigning examinees to modules targeted to different proficiency levels. Consequently, interactions between examinee groups and modules can create additional challenges for item calibration (Cai et al., 2021). Recent work using an operational MST program similarly showed that calibrating new items with small samples remains an important practical problem and that calibration performance depends on the estimation and data-collection design employed (Guo et al., 2024).

Machine-learning approaches offer an alternative way to address limitations of conventional estimation. Deep learning has increasingly been incorporated into psychometric modeling, including computationally efficient estimation for complex item factor models (Urban & Bauer, 2021). More recently, Belov et al. (2026) introduced a supervised neural network (NN) approach for estimating IRT parameters in small samples. Instead of estimating parameters directly from a likelihood function, the method derives features from simulated response data and trains an NN to map these features to known item parameters. Their simulation results showed that NN-based estimators could produce accurate parameter estimates with small samples and outperform Bayesian MCMC estimation under the conditions examined. Related work has also explored incorporating machine-learning-based item parameter predictions into small-sample calibration, emphasizing the potential value of prediction when calibration information is limited (Ulitzsch et al., 2026).

A potential limitation of supervised NN estimation, however, is its dependence on the conditions represented in the training data. Belov et al. (2026) identified the characteristics and distributions of simulated training data as an important area for further investigation. Although training and test distribution mismatch is not unique to adaptive testing, it may be relevant in MST because adaptive routing intentionally exposes different examinee groups to item sets

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 377–381
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

with different difficulty characteristics. An NN trained on one difficulty distribution may therefore be applied to operational data whose difficulty distribution differs because of modules, item-pool composition, or administration conditions. Evaluating this type of distributional shift provides evidence about whether NN-based parameter estimation generalizes beyond the conditions on which the model was trained.

2 Study Purpose

This study evaluates NN-based item difficulty estimation under MST conditions using relatively small Listening and Reading item banks. The simulation varies calibration sample size, the amount of simulated training information, and the degree of mismatch between the item difficulty distributions represented in the training and test data. This study addresses the following research questions:

RQ1. How do calibration sample size, item pool size, and the amount of simulated training information affect the accuracy of NN-based item difficulty estimates under MST conditions?

RQ2. To what extent does a mismatch between the item difficulty distributions represented in the training and test data affect the accuracy of NN-based Rasch item difficulty estimates under MST conditions?

3 Method

3.1 Data

The simulation design was based on operational characteristics of the WIDA ACCESS Listening and Reading assessments. Both domains use an MST design and are administered across five grade clusters: G01, G23, G45, G6-8, and G9-12. The Listening domain consists of six or eight stages, whereas Reading consists of eight or ten stages. The first two stages serve as routing stages. Based on performance during routing, examinees are assigned to Tier A, B, or C. Examinees assigned to Tiers B and C may progress to later stages that are not administered to students routed to a lower tier.

The number of items administered therefore differs according to the student's route. Listening tests contain either 18 or 24 administered items, and Reading tests contain either 24 or 30 items. Item difficulty parameters for both domains were estimated using the Rasch model and vertically scaled across grade clusters. These operational item parameter estimates served as the basis for constructing the simulation item pools. Table 1 presents the descriptive statistics of item difficulty parameters for Listening and Reading domain.

Grade	Item	Mean	P25	P75	SD
Listening					
G01	54	-1.69	-2.40	-0.59	1.23
G23	54	-0.37	-1.78	0.99	1.71
G45	54	0.22	-0.70	1.06	1.12
G68	54	0.30	-0.55	1.01	1.10
G91	54	1.15	0.34	1.94	1.12
Reading					
G01	72	-1.40	-2.16	-0.37	1.31
G23	72	-0.62	-1.35	0.29	1.16
G45	72	0.11	-0.80	1.11	1.23
G68	72	1.15	0.14	2.24	1.34
G91	72	1.61	0.96	2.47	1.18

Table 1: Descriptive statistics of item difficulty parameters for Listening and Reading domains across grade clusters

3.2 Simulation Design

The simulation evaluated NN-based estimation of item difficulty parameters under combinations of four factors: item pool size, calibration sample size, amount of simulated training information, and degree of training-test distribution mismatch. The factor levels followed the operational item-bank characteristics and extended the simulation framework used in previous NN-based IRT research.

The item pool factor contained two levels corresponding to the two operational domains: P = 270 for Listening and P = 360 for Reading. Calibration sample size was varied across seven levels, S = 10, 50, 100, 150, 200, 250, and 300. This range was selected to represent extremely limited through moderate calibration samples. The amount of simulated training information was manipulated through the number of iterations used to construct the training data: I = 5, 10, 15, and 20. In this study, an iteration represents an additional simulated MST response dataset generated from the item pool and therefore increases the amount of response-pattern information available for training the NN.

The final factor represented the correspondence between the item difficulty distributions used for

NN training and evaluation. Three conditions were examined. In the matched condition, test item difficulties were generated from the same difficulty distribution represented in the training data. In the moderate-shift condition, the mean test-item difficulty was shifted by 0.5 logits while in the extreme-shift condition, the mean test-item difficulty was shifted by 1.0 logit.

3.3 Training and Test Data Generation

For each simulation condition, training and test datasets were generated using the same MST routing structure. In Listening, six routing items were administered first, followed by 12 additional items for Tier A or 18 additional items for Tiers B and C. Reading data followed the corresponding operational MST map and domain-specific test lengths. Items were sampled from the relevant operational pool so that simulated examinees encountered only the items associated with their assigned route rather than the complete bank.

For each simulated examinee, latent proficiency was drawn from a standard normal distribution, theta ~ N(0, 1). Binary responses were then generated under the Rasch model as a function of examinee proficiency and item difficulty. Response-derived variables and normalized scores used as NN inputs were constructed following the supervised-learning procedure described by Belov et al. (2025). Repeating this process according to the assigned iteration condition produced the training information used to fit the estimator.

Independent test datasets were generated with the same response model, proficiency distribution, routing structure, and test lengths. The only systematic difference was the item-difficulty distribution: it was unchanged in the matched condition and shifted upward by 0.5 or 1.0 logit in the moderate- and large-shift conditions, respectively. This design holds the remaining data-generating features constant. One hundred Monte Carlo replications were attempted for each condition.

3.4 Neural Network Estimation

The NN-based estimators were implemented in Python using the TensorFlow library (TensorFlow Developers, 2023) for model training and item parameter estimation. Hyperparameter settings followed those used in Belov et al. (2025). For training, the generating item difficulties were used as item labels, and each item's response row was concatenated with the sorted normalized raw scores of the examinee sample to form a feature vector. The trained NN was then applied to feature vectors from independent test datasets to obtain item difficulty estimates.

The primary distinction from conventional IRT calibration is that the trained network does not re-estimate a likelihood for each test dataset. Instead, it applies a previously learned mapping from response-derived characteristics to item difficulty. This makes training data coverage a substantive part of the estimation problem: if the response patterns associated with new items fall outside the range represented during training, prediction error may increase even when the underlying Rasch model is correctly specified.

Estimation accuracy was evaluated using root mean squared error (RMSE) between the generating item difficulties and the NN estimates. RMSE summarizes the magnitude of parameter-recovery error on the Rasch difficulty scale, with smaller values indicating more accurate recovery. For each condition, RMSE was computed across evaluated items and summarized across Monte Carlo replications.

4 Results

Figures 1 and 2 compare the matched, moderate-shift, and large-shift conditions. For clarity, only lowest and highest iteration conditions (I = 5 and I = 20), are displayed. As expected, the large-shift condition generally produced highest RMSE, particularly when the number of training iterations was low. In contrast, the moderate-shift condition with I = 20 followed the matched-condition results closely across most sample sizes in both Listening and Reading. This pattern suggests that additional training information can substantially reduce the effect of a moderate mismatch between the training and test item-difficulty distributions. However, the same degree of recovery was not evident under the large-shift condition, where RMSE remained higher even with I = 20.

Overall, the results indicate that NN-based item parameter estimation benefits from both larger calibration samples and greater amount of simulated training information. More importantly, the effect of distributional mismatch appears to depend on its magnitude: moderate shifts can be largely accommodated when sufficient training information is available, whereas larger shifts continue to reduce estimation accuracy.

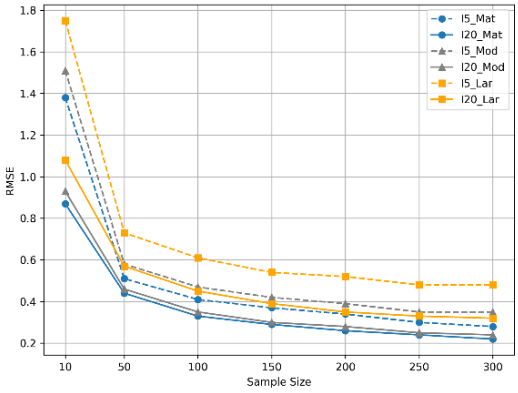


Figure 1: RMSE of item difficulty estimates for Listening across matched, moderate-shift, and large-shift conditions.

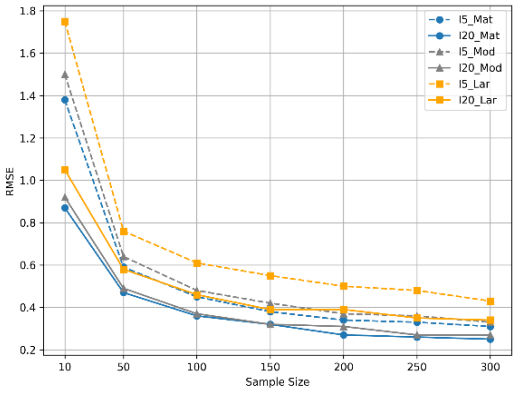


Figure 2: RMSE of item difficulty estimates for Reading across matched, moderate-shift, and large-shift conditions.

5 Conclusion

This study evaluated the performance and robustness of NN-based Rasch item difficulty estimation in an MST setting with relatively small operational item banks. Estimation accuracy improved as sample size and the amount of simulated training information increased, with the greatest gains occurring at smaller sample sizes and diminishing improvements at larger samples.

The findings also clarify the role of training–test distribution mismatch. Although estimation accuracy generally declined as the test-item difficulty distribution moved away from the training distribution, this effect was not uniform. Under the moderate-shift condition, estimates based on the higher number of training iterations closely approximated those obtained under the matched condition, suggesting that sufficient training information can provide meaningful robustness to moderate distributional changes. In contrast, a larger distribution shift continued to produce noticeably higher RMSE, even under the high iteration condition.

These findings suggest that NN-based IRT estimation may be practical for MST assessments with limited calibration samples, particularly when simulated training data are sufficiently rich and reasonably representative of operational item difficulty. The results also indicate that exact correspondence between training and operational distributions may not always be necessary; moderate differences may be tolerable when the NN is trained with adequate information. However, substantial distributional shifts remain an important limitation and should be considered when designing training simulations for operational use.

References

Ban, J. C., Hanson, B. A., Yi, Q., & Harris, D. J. (2002). Data sparseness and on-line pretest item calibration-scaling methods in CAT. *Journal of Educational Measurement*, *39*(3), 207–218. https://doi.org/10.1111/j.1745-3984.2002.tb01174.x

Belov, D. I., Lüdtke, O., & Ulitzsch, E. (2026). A supervised learning approach to estimating IRT models in small samples. *British Journal of Mathematical and Statistical Psychology, 79*(1), 66–94. https://doi.org/10.1111/bmsp.12396

Cai, L., Albano, A. D., & Roussos, L. A. (2021). An investigation of item calibration methods in multistage testing. *Measurement: Interdisciplinary Research and Perspectives, 19*(3), 163–178. https://doi.org/10.1080/15366367.2021.1878778

Guo, H., Johnson, M. S., McCaffrey, D. F., & Gu, L. (2024). Practical considerations in item calibration with small samples under multistage test design: A case study. *ETS Research Report Series, 2024*(1), 1–21. https://doi.org/10.1002/ets2.12376

TensorFlow Developers. (2023). *TensorFlow* (v2.21.0). Zenodo. https://zenodo.org/records/18894642

Ulitzsch, E., Belov, D., Lüdtke, O., & Robitzsch, A. (2026). Using item parameter predictions for reducing calibration sample requirements—A case study based on a high-stakes admission test. *Journal*

of Educational Measurement, *63*(1). https://doi.org/10.1111/jedm.12426

Urban, C. J., & Bauer, D. J. (2021). A deep learning algorithm for high-dimensional exploratory item factor analysis. *Psychometrika*, *86*(1), 1–29. https://doi.org/10.1007/s11336-021-09748-3

Validating LLM-Rated Item Features for Explanatory Item Response Models

Ahmed H. Bediwy
Pearson
ahmed.bediwy@pearson.com

Mubarak O. Mojoyinola
Ascend Learning
mubarak.mojoyinola@ascendlearning.com

Abstract

Explanatory item response models let test developers anticipate item difficulty from item design, but they depend on subject-matter experts rating every item on every hypothesized feature—a step that is slow, costly, and the practical bottleneck limiting how many items can be modeled. We ask whether large language models can supply those ratings. Four trained experts and six LLMs independently rated 15 grade-5 mathematics items on six psychometric features. We evaluate the LLM ratings twice: against the human consensus using quadratic-weighted κ and mixed-effects models, and against 1,452 student responses by using each source's ratings as the design matrix of a linear logistic test model (LLTM) benchmarked against a Rasch baseline. Claude Opus 4.7 performed best in recovering the item difficulty ($r = 0.89$), however, the human consensus against which agreement is measured performs worst at recovering Rasch difficulty ($r = 0.33$) compared to the rest of the LLMs. The reason is visible in the expert panel itself: inter-rater Fleiss' κ is at or below chance on three of six features, so the consensus is a noisy reference rather than ground truth. We also identify two failure modes (zero-variance and perfect collinearity) that agreement statistics cannot detect but make a feature unusable as an LLTM covariate. We argue that agreement should be reported alongside criterion-referenced recovery of item parameters, not in place of it.

1 Introduction

Modern assessment programs need far more items than expert committees can write and review. Automatic item generation addresses the writing bottleneck (Gierl and Lai, 2012; Morrison and Embretson, 2018), but it displaces the problem onto calibration: newly generated items have no response data, so their psychometric properties are unknown until they are field tested.

Explanatory item response theory offers a way out. Item-explanatory models such as the linear logistic test model (LLTM; Fischer, 1973) decompose Rasch item difficulty into the weighted contributions of theoretically motivated item features, so that difficulty can be anticipated from item design rather than recovered solely from responses (De Boeck and Wilson, 2004; Embretson, 1999). The approach has been applied to automatic item generation (Gierl and Lai, 2012), construct validity (Baghaei, 2008), item-position effects (Hohensinn et al., 2008; Demirkol and Kelecioğlu, 2022), and difficulty modeling (Hauenstein and Embretson, 2020; Toyama, 2021).

Every one of these applications requires the same expensive input. Before an LLTM can be estimated, subject-matter experts must rate every item on every hypothesized feature. This is the step that limits how many items can be examined and how quickly.

Large language models are an obvious candidate for that step, having matched or exceeded the performance of crowd and trained annotators on a range of text-labeling tasks (Gilardi et al., 2023). Work on predicting item difficulty directly from item text is also advancing (Benedetto et al., 2023; AlKhuzaey et al., 2024). What has not been examined is the intermediate case that matters for explanatory modeling: whether LLM-generated *features* yield explanatory IRT models that recover item parameters as well as human-rated features do.

Answering that question requires deciding what counts as validation. The default move is to treat expert ratings as ground truth and report agreement. We show that this default is misleading here. Our contributions are:

1. A dual evaluation of six LLMs as item-feature raters: agreement against a four-expert consensus, and criterion-referenced recovery of Rasch difficulty through an LLTM fitted to

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 382–388
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

1,452 student responses.

2. Evidence that the two criteria agree on how to rank *LLMs* but not on the standing of the *human reference*, which performs worst at recovering difficulty among all sources examined.
3. A diagnosis of why: the expert panel's own inter-rater reliability is at or below chance on three of six features, so agreement is measured against a noisy standard.
4. Two failure modes specific to LLM raters—zero-variance and perfect collinearity—that are invisible to agreement statistics but make the affected features unusable as LLTM covariates.

2 Method

2.1 Items and responses

The analysis uses 15 mathematics items for grade 5, administered as part of a formative assessment. The items measure three Common Core standards: 5.NBT.B.5, 5.NF.B.7, and 5.MD.B.2. Item difficulty is estimated from 1,452 student responses, which provide the empirical basis for the Rasch and LLTM models.

2.2 Features and rubric

Six features were rated on a three-point ordinal scale: Cognitive Load (CL), Visual Complexity (VC), Cognitive Domain (CD), Cognitive Coherence (CC), Response Variability (RV), and Standard Specificity (SS). Appendix A gives the operational definitions and anchors.

All six are potential candidate predictors of item difficulty in the LLTM models except for cases where a feature has zero-variance or perfect collinearity with another feature.

2.3 Rating panels

Four experienced subject-matter experts rated every item independently. The human consensus is the unique plurality of the four ratings; 2–2 splits were resolved by expert adjudication. Six LLMs–Claude Haiku 4.5 (Anthropic, 2025), Claude Opus 4.7 (Anthropic, 2026a), Claude Sonnet 4.6 (Anthropic, 2026b), GPT-5.2 (OpenAI, 2025), Gemini 3.1 Pro (Google DeepMind, 2026a), and Gemini 3.5 Flash (Google DeepMind, 2026b)–rated the same items under a common rubric and zero-shot prompting.

2.4 Aim 1: characterizing rating differences

Because the rubric scales are ordered, agreement between each LLM and the human consensus is quantified with quadratic-weighted Cohen's κ (Cohen, 1968), which penalizes a two-category disagreement more heavily than an adjacent one. Exact agreement, adjacent agreement, and ICC are also reported because weighted κ becomes unstable when one margin concentrates in a single category. Within-panel consistency is summarized by Fleiss' κ across the four experts and across the six LLMs separately.

Systematic differences are modeled with three linear mixed models (Bates et al., 2015), each with a random intercept for item, since ratings of the same item are dependent across raters:

$$\text{rating} \sim \text{panel} + (1 \mid \text{item}),$$

The human panel is the reference throughout, so a positive contrast means the LLM rates higher. Individual experts, not the derived consensus was used in these models, so the human side carries its own sampling variability.

2.5 Aim 2: recovery of item difficulty

The baseline is a Rasch model estimated by conditional maximum likelihood (Rasch, 1960; Mair and Hatzinger, 2007). Against it we fit an LLTM in which difficulty is constrained to a linear combination of the six rated features,

$$b_j = -\sum_{k=1}^{K} w_{jk}\,\eta_k, \tag{1}$$

where w_{jk} is the (standardized) rating of item j on feature k and η_k is the estimated feature weight. A separate LLTM is fitted per rating source—the human consensus, the AI consensus, and each LLM individually—so that the fits differ only in $\mathbf{W}$.

The quantity of interest is Pearson's correlation r between LLTM-reconstructed and freely estimated Rasch difficulties. We also report the likelihood-ratio test of the LLTM in comparison to Rasch baseline and RMSE.

Two problems arise with only 15 items and three rating categories, and both are properties of the rater rather than coding artifacts. A rating source that assigns every item the same category on a feature contributes a constant column(zero-variance); a rating source that rates two features identically contributes perfectly collinear columns.

Such columns are dropped during LLTM estimation, so the number of usable features K varies by rating source.

3 Results

3.1 LLM ratings agree only fairly with experts

Table 1 shows LLM agreement with the human consensus across the features. Mean quadratic-weighted κ separates the models clearly, from 0.50 for Claude Opus 4.7—moderate agreement—down to 0.10 for Claude Haiku 4.5. Only Opus exceeds the conventional threshold for moderate agreement; the remaining five sit in the fair range or below.

Agreement is far from uniform across features. Visual Complexity is rated consistently with experts by every model ($\kappa = 0.43$ to 0.78), whereas Response Variability is at or below chance for four of six models, reaching -0.49 for Gemini 3.1 Pro. Exact agreement (0.48–0.64) understates the picture somewhat, since adjacent agreement is high throughout (0.84–0.94): disagreements are usually one category, not two. With 15 items per model–feature cell these estimates are imprecise, and we read the pattern across features rather than any single coefficient.

3.2 Divergence is concentrated in two features

Pooled across features, the LLM panel rates only slightly higher than the expert panel ($\Delta = 0.12$, $SE = 0.05$, $p = .031$), and the item random intercept absorbs little variance (ICC $= 0.05$). After Bonferroni correction, no individual model differs significantly from the expert panel.

The interaction model locates the difference precisely (Table 2, Figure 1a). Four features show no detectable panel difference. Two show large ones: LLMs rate Cognitive Coherence 0.53 higher and Standard Specificity 0.54 higher than experts do (both $p < .001$). On a three-point scale, these are half-category shifts, and they push the LLM panel against the ceiling.

3.3 The expert panel is an unreliable reference

Table 3 and Figure 1b report within-panel consistency, and they reframe everything above. Fleiss' κ among the four experts reaches 0.15 on no feature at all, and is *negative* on three of the six: Cognitive Coherence (-0.08), Response Variability (-0.17), and Standard Specificity (-0.09). On those three features, the experts disagree with each other slightly more than chance would predict.

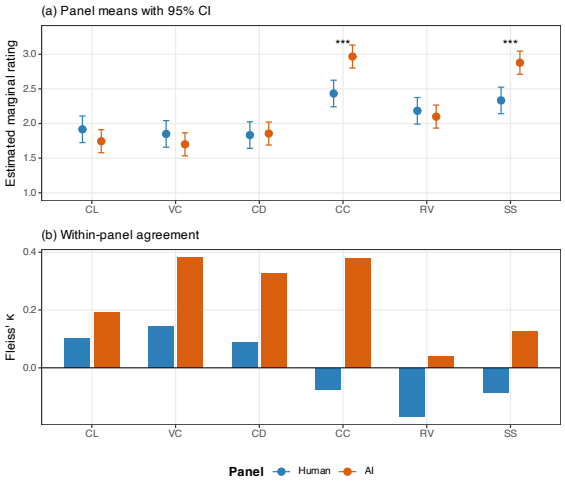


Figure 1: (a) Estimated marginal ratings by panel with 95% CIs; $^{***}p < .001$. (b) Within-panel agreement (Fleiss' κ). The LLM panel is more internally consistent than the expert panel on every feature.

The LLM panel is more consistent than the expert panel on every feature, in several cases by a wide margin (VC 0.38 versus 0.15; CC 0.38 versus -0.08). This does not make the LLMs correct—high agreement among models that share training data and prompt format is weak evidence—but it does mean that a low human–AI κ cannot be read as LLM error. Much of the disagreement quantified in Section 3.1 originates on the human side.

3.4 LLM features recover item difficulty better than expert features

The Rasch baseline fits well ($\log L = -8945.1$, 14 item parameters) and its difficulties correlate perfectly in rank with proportion correct ($\rho = -1.00$), confirming the calibration.

Table 4 reports the LLTM results. The ranking of the six LLMs based on how well they recover the rasch difficulty is exactly the same as the agreement ranking produced in Section 3.1: Opus, GPT-5.2, Sonnet, Flash, Gemini 3.1 Pro, and Haiku, in that order for both criteria. Agreement with experts, for all its limitations, does rank these models the same way their downstream psychometric model parameter recovery does.

The human consensus is where the two criteria differ. It is the *worst* of all eight sources used to estimate LLTMs: its six-feature LLTM recovers Rasch difficulty at $r = 0.33$ ($R^2 = 0.11$). Five of six LLMs do better with the same or fewer covariates.

Model	CL	VC	CD	CC	RV	SS	Mean
Claude Opus 4.7	0.43	0.78	0.52	—	-0.18	0.94	0.50
GPT-5.2	0.09	0.76	0.08	0.36	0.06	0.55	0.32
Claude Sonnet 4.6	0.45	0.58	0.28	0.36	-0.12	0.23	0.30
Gemini 3.5 Flash	0.30	0.64	0.16	—	-0.08	—	0.26
Gemini 3.1 Pro	0.04	0.64	0.40	0.36	-0.49	0.23	0.20
Claude Haiku 4.5	0.29	0.43	-0.40	—	0.08	—	0.10

Table 1: Quadratic-weighted κ between each LLM and the human consensus, by feature, ordered by the mean. Dashes mark cells where κ is undefined because the model assigned a single category to all 15 items. Feature codes are defined in Section 2.2.

Feature	Human	AI	Δ	SE	t	p	
CL	1.917	1.744	-0.172	0.110	-1.567	0.117	
VC	1.850	1.700	-0.150	0.110	-1.365	0.173	
CD	1.833	1.856	0.022	0.110	0.202	0.840	
CC	2.433	2.967	0.533	0.110	4.853	0.000	***
RV	2.183	2.100	-0.083	0.110	-0.758	0.448	
SS	2.333	2.878	0.544	0.110	4.954	0.000	***

Table 2: Estimated marginal ratings by panel and the AI − Human contrast from the panel × feature mixed model (Bonferroni-adjusted). $^{***}p < .001$.

Feature	κ_F Human	κ_F AI	ICC Human	ICC AI
CL	0.101	0.193	0.407	0.337
VC	0.145	0.382	0.632	0.880
CD	0.089	0.328	0.381	0.803
CC	-0.076	0.379	0.255	0.800
RV	-0.169	0.040	-0.159	0.627
SS	-0.085	0.128	0.391	0.596

Table 3: Within-panel inter-rater agreement: Fleiss' κ and average-measure ICC for the four experts and the six LLMs.

Claude Opus 4.7 reaches $r = 0.89$ ($R^2 = 0.79$) using five features; GPT-5.2 reaches $r = 0.85$ and Claude Sonnet 4.6 $r = 0.81$, both with six. The gap is not a parameter-count artifact: the human consensus has the *most* covariates available and still performs worst.

Figure 2 makes the difference concrete. Under the human consensus, the reconstructed difficulties are compressed into a narrow band around zero–the features carry almost no signal about which items are hard. Under Claude Opus 4.7, they track the Rasch estimates closely across the full range.

The likelihood-ratio tests all reject the LLTMs ($p < .001$ throughout), as expected at this sample size.

3.5 Zero-variance and perfect collinearity

Two failure modes emerge that agreement statistics cannot detect.

Zero-variance. The AI consensus assigns category 3 to all 15 items on both Cognitive Coherence and Standard Specificity. The variance is exactly zero, so neither feature can be used in the LLTM: the AI consensus is left with four usable covariates rather than six. The same occurs for Claude Haiku 4.5 and Gemini 3.5 Flash, and for Cognitive Coherence alone in Claude Opus 4.7. This is the Aim 1 result seen from another angle–the two features on which LLMs significantly overrate are the two they push to the ceiling.

Perfect collinearity. Gemini 3.1 Pro assigns identical ratings to Cognitive Coherence and Standard Specificity on every item ($r = 1.00$), making the columns perfectly collinear. Therefore, one of them has to be drop during model estimation.

Table 5 shows the consequence for the estimated weights. Where the human weights are uniformly small ($|\eta| \leq 0.18$), the AI consensus places large weights on Visual Complexity (0.52) and Cognitive Domain (-0.62) which conforms to difficulty theory, whereas the noisy human consensus does not.

Figure 3 shows that the sources disagree about which features drive difficulty. The largest weight belongs to a different feature for almost every source: Cognitive Load for Claude Opus 4.7 and Gemini 3.5 Flash, Visual Complexity for GPT-5.2, Cognitive Domain for the human and AI consensus, Cognitive Coherence for Claude Sonnet 4.6 and Gemini 3.1 Pro, and Response Variability for Claude Haiku 4.5. Recovering difficulty well and recovering it *for the same reasons* are different achievements, and only the first is established here.

4 Discussion

The two validation criteria agree about the LLMs and disagree about the humans, and that asymmetry is the finding. Ranking the six models by agreement with experts reproduces their ranking by dif-

Rating source	K	$\log L$	BIC	LR	df	r	R^2	RMSE
Rasch (baseline)	—	-8945.11	17992.14	—	—	1.00	1.00	0.00
Claude Opus 4.7	5	-9279.83	18596.06	669.44	9	0.89	0.79	0.41
GPT-5.2	6	-9391.41	18826.51	892.61	8	0.85	0.72	0.47
Claude Sonnet 4.6	6	-9488.60	19020.89	1086.99	8	0.81	0.66	0.52
Gemini 3.5 Flash	4	-9703.44	19436.00	1516.67	10	0.70	0.49	0.64
Gemini 3.1 Pro	5	-9849.80	19736.01	1809.40	9	0.65	0.43	0.68
AI consensus	4	-10069.71	20168.54	2249.20	10	0.52	0.27	0.77
Claude Haiku 4.5	4	-14039.58	28139.09	2558.73	10	0.34	0.12	0.84
Human consensus	6	-10282.84	20609.36	2675.46	8	0.33	0.11	0.85

Table 4: LLTM fit and recovery of Rasch difficulty by rating source. K is the number of usable feature covariates after dropping constant and aliased columns; r is the correlation between LLTM-reconstructed and freely estimated Rasch difficulties. Sources are ordered by r.

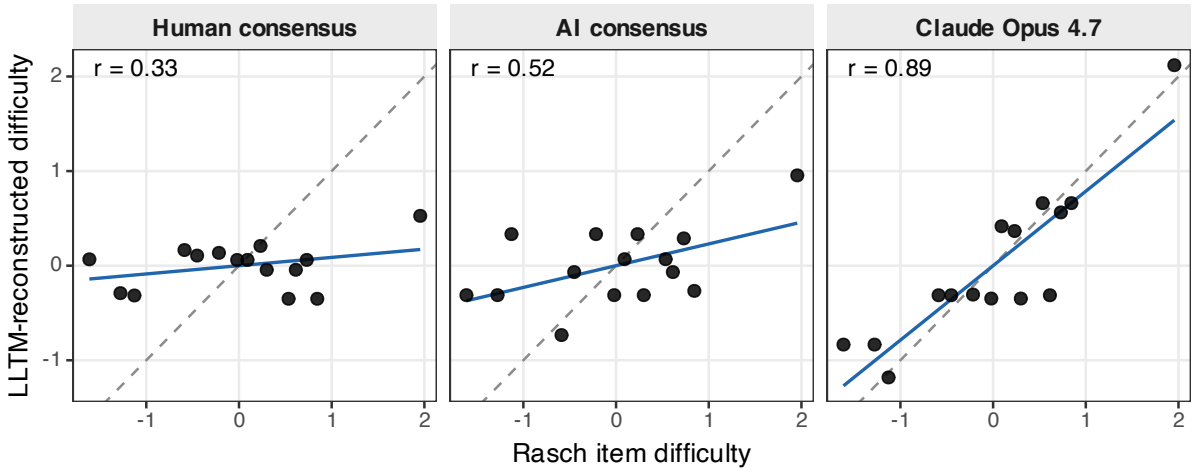


Figure 2: LLTM-reconstructed versus freely estimated Rasch difficulty for three rating sources. The dashed line is the identity line; the solid line is the least-squares fit.

Feature	Human consensus	AI consensus
CL	-0.067 (0.019)	-0.164 (0.028)
VC	0.176 (0.024)	0.518 (0.027)
CD	-0.184 (0.019)	-0.617 (0.035)
CC	-0.127 (0.023)	constant
RV	0.089 (0.018)	-0.197 (0.028)
SS	0.123 (0.025)	constant

Table 5: LLTM feature weights η (standard errors) on the difficulty scale, for standardized covariates. "constant" marks features the source rated identically on all 15 items.

ficulty recovery exactly. Agreement is therefore doing useful work: it is a cheap, response-free proxy that orders candidate raters correctly. What it cannot do is tell you where the reference standard itself sits on that scale–and here the expert consensus lands at the bottom, recovering difficulty at $r = 0.33$ against 0.89 for the best model.

The resolution is that the expert consensus is not ground truth. With Fleiss' κ at or below zero on three features, aggregating four experts produces a reference standard that is substantially noisy. An agreement statistic computed against it is bounded above by its reliability, so "only fair agreement with experts" and "better than experts at explaining difficulty" are not in tension once the reference is recognized as noisy.

This has a practical consequence for anyone considering LLMs for item review. Agreement with a small expert panel is cheap and is the field's default, and our results support using it to *rank* candidate raters. They do not support using it to decide whether a rater is good enough, because that judgment is only as good as the panel. Where response data exist, criterion-referenced recovery of item parameters answers the question agreement cannot, and the two should be reported together.

The zero-variance and perfect collinearity problems identified in some LLM ratings sharpen the recommendation. Both are silent in agreement statistics–a model that rates every item 3 can still post a respectable exact-agreement rate when the

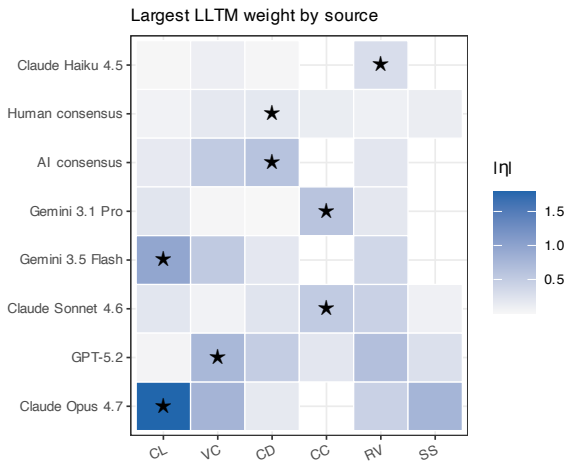


Figure 3: Absolute LLTM feature weights by rating source.

human mode is also 3, but both are fatal for explanatory modeling, since a constant or collinear covariate carries no information. Screening candidate raters for category variance and inter-feature collinearity costs nothing and should precede any agreement analysis.

These results support a narrower claim than "LLMs can replace expert raters." Three features concentrate most of the trouble—Cognitive Coherence, Response Variability, and Standard Specificity. They are where the experts disagree with each other, where the LLMs saturate, and where one model collapses two rubric entries onto a single judgment. Yet several models place their largest difficulty weights on exactly these features. A feature can be unreliable, saturated, and still carry signal, which points at the rubric rather than the raters: these three anchors need sharper operational definitions before either panel can be held to them, and that is where we would direct effort next.

Limitations

The item sample is the binding constraint. Fifteen items with response data support a Rasch calibration comfortably but leave the LLTM comparison badly underpowered. The ordering of sources in Table 4 is consistent and large, but this study cannot establish its magnitude, and we do not claim statistical support for the difference. Extending the calibration sample is the single most valuable next step.

The human panel consists of four raters. Their low inter-rater reliability is central to our interpretation, yet it is itself estimated imprecisely from 15 items, and we cannot separate genuine construct ambiguity from insufficient rater training. A larger and better-calibrated expert panel might close the gap we report.

Each LLM rated each item once, using default settings and a single prompt. We therefore cannot distinguish stable model behavior from prompt sensitivity or sampling variability, and the within-LLM agreement we report is inflated relative to what independent replications would show. The models also share architectural lineage and training data, so their mutual agreement is not evidence of validity.

Recovery of Rasch difficulty is a proxy for usefulness, not a direct measure of it. A design matrix could recover in-sample difficulty well and still fail to predict parameters for genuinely new items, which is the use case that motivates explanatory modeling. We report in-sample recovery only; out-of-sample prediction for generated items is left to future work.

Finally, the findings are specific to grade-5 mathematics, three Common Core standards, and this six-feature rubric. Whether the dissociation between agreement and recovery generalizes to other domains, grade levels, or feature sets is untested.

References

Samah AlKhuzaey, Floriana Grasso, Terry R. Payne, and Valentina Tamma. 2024. Text-based question difficulty prediction: A systematic review of automatic approaches. *International Journal of Artificial Intelligence in Education*, 34(3):862–914.

Anthropic. 2025. Claude haiku 4.5. Large language model.

Anthropic. 2026a. Claude opus 4.7. Large language model. Model identifier: claude-opus-4-7.

Anthropic. 2026b. Claude sonnet 4.6. Large language model. Model identifier: claude-sonnet-4-6.

Purya Baghaei. 2008. The Rasch model as a construct validation tool. *Rasch Measurement Transactions*, 22(1):1145–1146.

Douglas Bates, Martin Mächler, Ben Bolker, and Steve Walker. 2015. Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1):1–48.

Luca Benedetto, Paolo Cremonesi, Andrew Caines, Paula Buttery, Andrea Cappelli, Andrea Giussani, and Roberto Turrin. 2023. A survey on recent approaches to question difficulty estimation from text. *ACM Computing Surveys*, 55(9):1–37.

Jacob Cohen. 1968. Weighted kappa: Nominal scale agreement provision for scaled disagreement or partial credit. *Psychological Bulletin*, 70(4):213–220.

Paul De Boeck and Mark Wilson, editors. 2004. *Explanatory Item Response Models: A Generalized Linear and Nonlinear Approach*. Springer, New York.

Sinem Demirkol and Hülya Kelecioğlu. 2022. Analyzing the interaction of item position effect and student characteristics within explanatory IRT models. *Journal of Measurement and Evaluation in Education and Psychology*, 13(4):282–304.

Susan E. Embretson. 1999. Generating items during testing: Psychometric issues and models. *Psychometrika*, 64(4):407–433.

Gerhard H. Fischer. 1973. The linear logistic test model as an instrument in educational research. *Acta Psychologica*, 37(6):359–374.

Mark J. Gierl and Hollis Lai. 2012. The role of item models in automatic item generation. *International Journal of Testing*, 12(3):273–298.

Fabrizio Gilardi, Meysam Alizadeh, and Maël Kubli. 2023. ChatGPT outperforms crowd workers for text-annotation tasks. *Proceedings of the National Academy of Sciences*, 120(30):e2305016120.

Google DeepMind. 2026a. Gemini 3.1 pro. Large language model.

Google DeepMind. 2026b. Gemini 3.5 flash. Large language model.

Clifford E. Hauenstein and Susan E. Embretson. 2020. Modeling item difficulty in a dynamic test. *Journal of Cognitive Education and Psychology*, 19(2):93–106.

Christine Hohensinn, Klaus D. Kubinger, Manuel Reif, Stefana Holocher-Ertl, Lale Khorramdel, and Martina Frebort. 2008. Examining item-position effects in large-scale assessment using the linear logistic test model. *Psychology Science Quarterly*, 50(3):391–402.

Patrick Mair and Reinhold Hatzinger. 2007. Extended Rasch Modeling: The eRm package for the application of IRT models in R. *Journal of Statistical Software*, 20(9):1–20.

Kristin M. Morrison and Susan E. Embretson. 2018. Item generation. In Paul Irwing, Tom Booth, and David J. Hughes, editors, *The Wiley Handbook of Psychometric Testing: A Multidisciplinary Reference on Survey, Scale and Test Development*, volume 1, pages 75–94. Wiley.

OpenAI. 2025. GPT-5.2. Large language model.

Georg Rasch. 1960. *Probabilistic Models for Some Intelligence and Attainment Tests*. Danish Institute for Educational Research, Copenhagen.

Yukie Toyama. 2021. What makes reading difficult? an investigation of the contributions of passage, task, and reader characteristics on comprehension performance. *Reading Research Quarterly*, 56(4):633–642.

A Rating Rubric

All raters, human and LLM, received the same six-feature rubric. Each feature is rated 1–3 with the anchors below; the full prompt is released with the analysis code.

Cognitive Load Number of mental steps required. 1 = single-step retrieval or one operation; 2 = two connected operations, or one operation plus context interpretation; 3 = three or more operations or nested reasoning.

Visual Complexity Density of visual information. 1 = text only or a simple table; 2 = a clear diagram or basic graph that supports understanding; 3 = a complex or cluttered visual requiring interpretation.

Cognitive Domain Predominant type of mathematical thinking. 1 = procedural (applying an algorithm); 2 = conceptual (understanding relationships or why a procedure works); 3 = application (solving a contextual problem strategically).

Cognitive Coherence Clarity and singularity of the solution path. 1 = multiple valid pathways or an ambiguous statement; 2 = one clear intended path with possible alternatives; 3 = a single obvious strategy and an unambiguous statement.

Response Variability Quality of the distractors. 1 = arbitrary distractors inviting random guessing; 2 = distractors representing some common errors; 3 = distractors systematically representing documented misconceptions.

Standard Specificity Precision of alignment to one standard. 1 = draws on several standards or grade levels; 2 = one main standard plus supporting knowledge from another; 3 = tightly focused on a single standard.

Efficacy of Student–AI Co-Authored Math Word Problems in an Intelligent Tutoring System

Kole Norberg, April Murphy, Stephen Fancsali, and Steve Ritter
Carnegie Learning
436 Seventh Avenue, Suite 200
Pittsburgh, PA 15219
knorberg1@carnegielearning.com

Abstract

We evaluated student ($N = 1{,}361$) interest and performance on student-AI co-authored math word problems. Performance matched or exceeded standard problems. Students rated peer-authored problems more often, especially when authorship was disclosed. Liking predicted first-attempt accuracy when problems required greater textual engagement, supporting interest-based context personalization.

1 Background

Mathematics story problems require students to coordinate mathematical reasoning with the interpretation of narrative text (Koedinger and Nathan, 2004). As a result, contextual familiarity with the narrative can influence reading comprehension, cognitive load, and motivation and thereby have an effect on student performance (Walkington et al., 2015). A student unfamiliar with the vocabulary or cultural assumptions in a problem may face demands unrelated to the intended mathematical construct, so performance differences across problems may partly reflect access to context rather than mathematical knowledge (Walkington et al., 2018).

Context personalization aims to reduce these construct-irrelevant demands, and aligning problems with student interests has been found to improve accuracy, speed, and engagement (Bernacki and Walkington, 2018; Clinton and Walkington, 2019; Walkington, 2013). These benefits appear inconsistently, however, sometimes producing small or null effects (Fancsali and Ritter, 2014) and sometimes distracting students when extraneous details fail to integrate with the mathematics (Harp and Mayer, 1998; Magner et al., 2014). In addition to requiring the appropriate grain size of personalization, personalizing problem context for each student can present a challenge due to students' wide variety of interests (Walkington, 2026; Walkington and Hayata, 2017).

Generative AI offers a means of producing a broader, more responsive range of contexts. The AI Math Personalization Tool (AMPT) supported students in co-authoring story problems around their interests while preserving the mathematical structure required by the curriculum (Norberg et al., 2026). After writing problems with AMPT and learning their problems would be added to the Intelligent Tutoring System they used in their math class (MATHia), students' sense of belonging in mathematics improved (Norberg et al., 2025b). Although the authoring approach succeeded as a method to both engage students and generate larger and more varied samples of problem contexts, it was not clear whether these student–AI co-authored problems would pose additional difficulties for students who did not author them. Language models are prone to writing longer texts with less familiar vocabulary (Norberg et al., 2023) both of which can decrease performance in mathematics (Abedi and Lord, 2001; Norberg et al., 2025a,c; Walkington et al., 2018). Further, some contexts preferred by students while co-authoring the problems might not be familiar or interesting to a broader cohort. Therefore, in this study we tested if the student-AI authored problems created any additional challenges to students over standard problems.

A further question concerned authorship itself. Co-authoring problems with AMPT raised students' sense of belonging in mathematics (Norberg et al., 2025b), and students who later solve those problems might also respond to knowing that a classmate wrote them. Thus, we tested how notifications of authorship affected students' ratings and performance on the problems.

2 Aims

We examined a large-scale deployment of student co-authored problems to ask (a) whether students who did not write a problem perform as well on it

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 389–396
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

as on standard problems, (b) whether telling students that a peer authored the problem changes their performance and how they rate the problem context, (c) whether students like the same problem context, and (d) whether liking a problem context predicts performance on it.

3 Sample

Middle school students (initial n = 1,452) working through MATHia (an Intelligent Tutoring System) as standard educational practice were randomly assigned using UpGrade (Ritter et al., 2020) to three conditions: Co-Authored, Co-Authored + Notification, or Control. Six percent of students (91) assigned to a co-authored condition never encountered a co-authored problem and were dropped. The dataset comprised 2,049 problems (63% co-authored) across seven lessons on the topics of probability, algebra, and ratios. Lessons within a topic share a format, so we report results by topic. Random assignment occurred within topic which resulted in some students ($n = 141$) participating in more than one condition across separate topics. Across conditions, there were 525 control, 497 co-authored, and 480 co-authored-plus-notification student-condition enrollments.

4 Method

In the co-authored conditions, students saw a mixture of co-authored and standard problems; in the notification condition, co-authored problems carried the message "This problem was authored by a MATHia student"; control students received no co-authored problems. In both co-authored conditions, students could rate problems with a thumbs up or down (Figure 1). Because we hypothesized that rating alone may affect students' experience, control was kept as a true business as usual condition and did not provide an option for rating.

In MATHia, a problem contains multiple steps, each associated with a skill the student is working to master. Students cannot move to the next problem until every step is answered correctly, so a student may make several errors on one step. MATHia traces each student's mastery of a skill with Bayesian knowledge tracing (Corbett and Anderson, 1995; Ritter et al., 2007). Once all skills in a lesson are mastered, the student is said to have mastered the lesson. In this study, if a student reached 25 problems without mastering all of the skills the lesson concluded but with a flag noting that the lesson was not mastered. Problems within a subtopic differ in difficulty according to the skills they assess, so each co-authored scenario was spread across problems assessing different skills following the same procedure for creating standard MATHia problems. For example, two problems could use the same scenario but assess different relations or skills.

All lessons utilized stories as part of the context; however, in the algebra topic the standard MATHia problems follow a highly predictable structure in which the problem can be answered by searching the text for the numbers and a few words which explain the relationship. In the ratio problems, on the other hand, the structure of the text is more varied. See Table 1 for descriptive statistics concerning the texts in each topic. We considered this variation in our models to understand if effects may vary based on how similar the texts within a topic are to each other.

Ratings asked "Do you like this problem's storyline" and provided a thumbs up and thumbs down option. Ratings were optional and could be submitted before or after solving. 84.2% were submitted before solving. Although we asked students to rate the storyline, students might also be influenced by their performance on the problem. For this reason, we restricted analysis involving ratings to those provided before the student attempted to solve any steps in the problem and further controlled for problem difficulty. Problem difficulty was measured as each problem's error rate across all students who have encountered it in MATHia, defined independently of this study's conditions.

Lesson mastery was modeled among all student-lessons. Time to mastery, error rates, and first-attempt correctness were modeled only among the students who graduated the lesson, to avoid conflating these results with the graduation model. Students could complete more than one lesson. Therefore, lesson-level models used crossed random intercepts for students (nested within school) and lesson. Problem-level models used crossed random intercepts for students (nested within school) and problems (nested within lesson). Log odds are reported in tables but back transformed to odds ratios within the text.

Topic type	Words	Sentence length	Reading ease	Hapax ratio	Vocabulary variance
Picture algebra	66.9 / 81.6	10.3 / 10.3	68.7 / 72.8	0.13 / 0.20	0.62 / 0.80
Ratio and proportion	67.0 / 78.5	13.5 / 14.7	76.4 / 73.9	0.16 / 0.25	0.86 / 0.86
Probability	47.4 / 61.2	12.7 / 13.8	83.3 / 75.1	0.37 / 0.44	0.87 / 0.89

Table 1: Text Measures by Topic Type

Note. Each cell gives standard / peer-authored values. Words is words per problem, sentence length is tokens per sentence, and reading ease is Flesch reading ease. Hapax ratio is the proportion of a problem's words that appear only once. Vocabulary variance is the mean pairwise distance between problems' content-word sets within a cell, where 1 indicates that two problems share no content words.

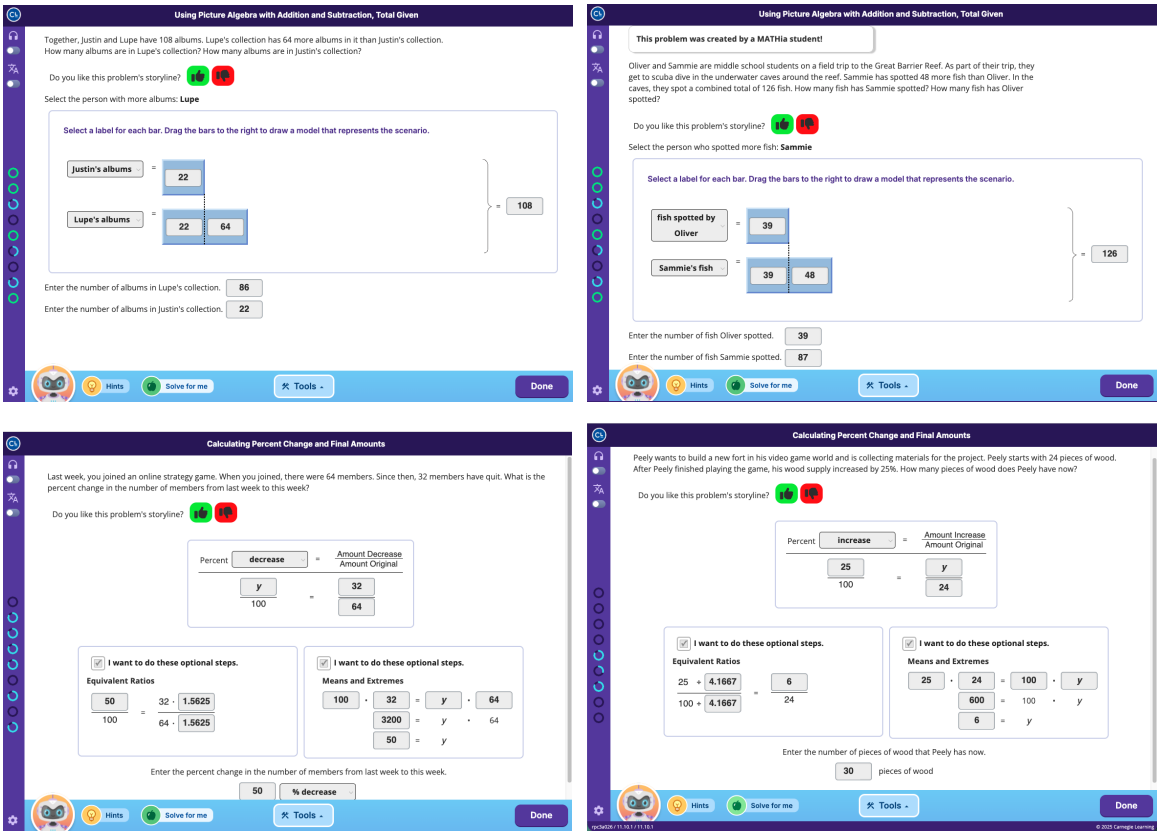


Figure 1: Standard (left) and peer-authored (right) picture-algebra problems with notification (top) and ratio-and-proportion problems without notification (bottom).

5 Results

5.1 Lesson-level outcomes

Lesson mastery rates in the co-authored conditions did not reliably differ from control nor did errors per step among students who mastered the lesson (all p-values $> .45$). However, students in the Co-Authored + Notification condition took 2.68 min longer than students in the Control condition to master a topic, 95% CI [0.38, 4.98], $t(1851.3) = 2.28$, $p = .023$. Students in the Co-Authored condition did not significantly differ in time to mastery from students in the Control condition, $p = .293$ (Table 2).

5.2 First attempt success at a step

The prior outcomes reflect broad lesson-level effects, but students in the Co-Authoring and Co-Authoring + Notification conditions saw a mixture of peer-authored and standard problems. In this section we consider how often students answered a step correctly on their first attempt for standard as compared to peer-authored problems. The model contains a term for problem type and a term for condition. The intercept is a standard problem in the Control condition, the condition terms give standard problems in each co-authoring condition relative to that intercept, and the problem-type term gives peer-authored problems relative to the standard problems in the same condition.

The problem-type term compares the two kinds of problems a student saw in the same lesson. Students had 1.12 times greater odds of answering a step correctly on their first attempt for peer-authored problems than for standard problems, 95%

Predictor	B	SE	z / t	p
Mastery (logistic)				
Intercept	3.74	0.59	6.38	$< .001$
Co-Authoring	0.22	0.30	0.74	.459
Co-Authoring + Notification	-0.04	0.29	-0.15	.877
Lesson-pooled error rate (linear; errors per step)				
Intercept	0.19	0.02	12.34	$< .001$
Co-Authoring	0.004	0.01	0.53	.598
Co-Authoring + Notification	0.005	0.01	0.74	.459
Time to mastery (linear; minutes)				
Intercept	29.01	4.78	6.07	$< .001$
Co-Authoring	1.22	1.16	1.05	.293
Co-Authoring + Notification	2.68	1.18	2.28	.023

Table 2: Mixed-Model Estimates for Lesson-Level Mastery, Time to Mastery, and Pooled Error Rate

Note. Reference = Control. Mastery uses all student-lessons. Time and pooled-error models include mastered lessons only. Wald z and Satterthwaite t tests are reported.

CI [1.07, 1.18], $z = 4.45$, $p < .001$.

The condition terms compare performance on standard problems across conditions. Students in the Co-Authoring condition had slightly lower odds of answering correctly than students in the Control condition, $OR = 0.94$, 95% CI [0.88, 1.00], $z = -2.00$, $p = .046$. Performance did not significantly differ between the Co-Authoring + Notification and Control conditions, $p = .213$. See Table 3.

5.3 Ratings

Due to a technical error, ratings for the probability topic were not recorded and cannot be analyzed. Because the control condition contained no rating opportunity, the models in this section use only the two co-authoring conditions. To control for differences in rating and liking based on difficulty, each model includes problem difficulty, the problem's error rate across all students who have encountered it in MATHia.

5.3.1 Rating frequency

Students rated 45.3% of eligible problems, and they rated peer-authored problems more often than standard problems in both conditions, with a larger difference when the notification was present. With the notification present, students rated 50.4% of peer-authored problems and 42.8% of standard problems, $OR = 1.81$, 95% CI [1.51, 2.16], $z = 6.47$, $p < .001$. Without it, students rated 45.7% of peer-authored problems and 41.6% of standard problems, $OR = 1.26$, 95% CI [1.07, 1.49], $z = 2.72$, $p = .007$. The difference was 1.43 times larger with the notification, 95% CI [1.13, 1.83], $z = 2.94$, $p = .003$ (Table 4).

Students also rated harder problems more often. Each increase of 0.10 in a problem's error rate raised the odds of a rating by a factor of 1.11, 95% CI [1.04, 1.18], $z = 3.15$, $p = .002$.

5.3.2 Rating across the lesson

Students rated problems more often early in the lesson and tapered off as the lesson went on, but a notification of peer-authorship reduced the slope. Among peer-authored problems, the odds that a student rated a problem fell by a factor of 0.89 with each successive problem in the Co-Authoring condition, 95% CI [0.87, 0.92], $z = -6.98$, $p < .001$, a decline from 55.7% on the first problem to 33.9% by the tenth. Under the notification the odds fell by a factor of 0.96, 95% CI [0.93, 0.98], $z = -3.08$, $p = .002$, a decline from 52.5% to 43.9%. The two slopes differed, $OR = 1.07$, 95% CI [1.03, 1.12], $z = 3.08$, $p = .002$, while the slopes for standard problems did not significantly differ across the two conditions, $p = .655$ (Figure 2; Table 4).

5.3.3 Liking

Students gave a thumbs up to 59.3% of the problems they rated. They did not significantly favor peer-authored problems over standard problems, $p = .238$. Students in the two conditions also did not significantly differ in how positively they rated problems, $p = .850$. Liking did not significantly vary with problem difficulty, $p = .387$. Ratings did grow more positive as a lesson went on, $OR = 1.03$ per successive problem, 95% CI [1.00, 1.05], $z = 2.35$, $p = .019$.

Predictor	B	SE	z	p
Intercept	1.77	0.12	15.13	$< .001$
Problem-Type (Peer-Authored)	0.12	0.03	4.45	$< .001$
Co-Authoring Condition	-0.07	0.03	-2.00	.046
Co-Authoring + Notification Condition	-0.04	0.03	-1.24	.213

Table 3: First-Attempt Correctness

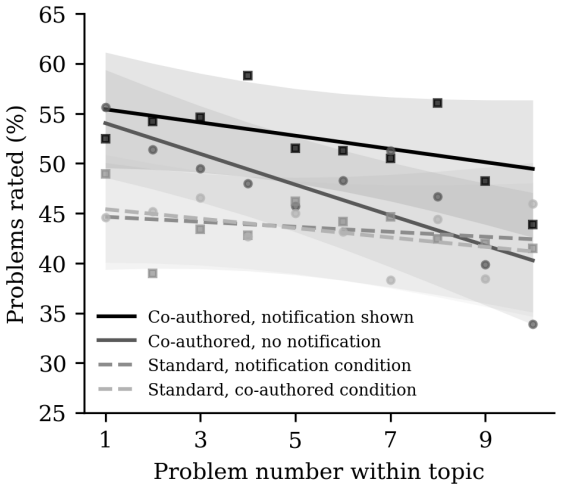


Figure 2: Percentage of problems rated by problem number within a lesson, for peer-authored and standard problems in each co-authoring condition, with observed points and logistic fits.

A null logistic model attributed 53.5% of the variance in liking to students (intercept variance = 3.79) and almost none to scenarios (0.009; 0.1%; Table 4). Whether a problem was liked depended overwhelmingly on who was rating, not on which scenario appeared.

5.4 Liking predicts success

A thumbs up submitted before solving predicted higher first-attempt success even when controlling for problem difficulty, but only in topic areas where the text of the problem varied more across problems (Figure 3; Table 5). In the ratio and proportion topics, students answered a step correctly on the first attempt at 1.20 times the odds when they had rated the problem positively, 95% CI [1.03, 1.40], $z = 2.36$, $p = .018$, a difference of 4.7 points in observed accuracy, 84.4% against 79.7%. In picture algebra, where the text is less varied across problems, liking did not predict first-attempt accuracy, $p = .055$, and observed accuracy was 89.3% and 89.4%. The interaction of liking with topic type was reliable, $OR = 1.33$, 95% CI [1.11, 1.60], $z = 3.05$, $p = .002$, in a model that also contained problem difficulty, problem type, problem position,

Predictor	B	SE	z	p
Rating a problem				
Intercept	-1.08	0.48	-2.26	.024
Problem Type	0.23	0.09	2.72	.007
Notification	-0.31	0.17	-1.77	.077
Problem Number	-0.07	0.01	-8.81	$< .001$
Problem Difficulty	1.04	0.33	3.15	.002
Problem Type × Notification	0.36	0.12	2.94	.003
Rating a peer-authored problem across the lesson				
Intercept	-0.58	0.48	-1.19	.233
Problem Number	-0.11	0.02	-6.98	$< .001$
Notification	-0.22	0.23	-0.97	.332
Problem Difficulty	0.69	0.38	1.83	.067
Problem Number × Notification	0.07	0.02	3.08	.002
A rating being positive				
Intercept	0.87	0.33	2.61	.009
Problem Number	0.03	0.01	2.35	.019
Problem Type	-0.10	0.08	-1.18	.238
Notification	0.03	0.18	0.19	.850
Problem Difficulty	-0.39	0.46	-0.87	.387

Variance components for liking	**Variance**	**ICC**
Student	3.79	.54
Scenario	0.01	.001
Residual (logistic)	3.29	.46

Table 4: Rating and Liking Across a Lesson

Note. Reference = Co-Authoring without notification. The first model uses all problem encounters ($n = 10{,}021$), the second peer-authored encounters ($n = 5{,}325$), and the third the ratings students submitted ($n = 4{,}540$). The variance components come from a null model over the same ratings with random intercepts for student and scenario where the residual is fixed at $\pi^2/3 = 3.29$.

and condition. Structural differences between the topic types may explain this interaction. Ratio and proportion problems are longer and share few content words with one another, while picture algebra problems restate one relation in a fixed frame and share extensive vocabulary, with a text variety of 0.62 against 0.86 for ratio and proportion (Table 1). The peer-authored problems introduce new variance in vocabulary to the picture algebra problems which may have reduced any effect of liking on performance.

6 Discussion

Context personalization has been studied as something done for a student, with the system selecting or generating a setting to match that student's interests (Bernacki and Walkington, 2018; Clinton and

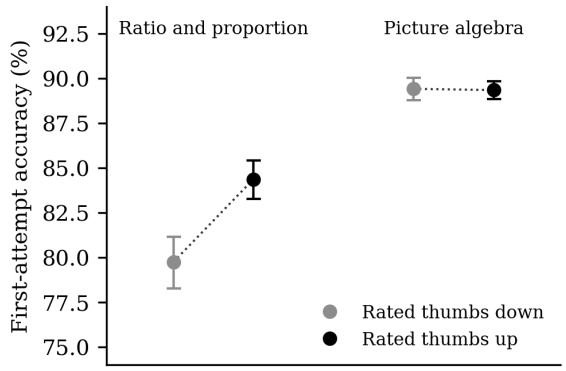


Figure 3: First-attempt accuracy on steps of problems rated before solving, by rating and topic type, with 95% Wilson intervals.

Predictor	B	SE	z	p
Intercept	2.33	0.21	11.28	$< .001$
Thumbs Up	-0.10	0.05	-1.92	.055
Ratio and Proportion	-0.80	0.26	-3.07	.002
Peer Auth Problem	0.01	0.04	0.30	.762
Problem Number	0.09	0.01	15.83	$< .001$
Problem Difficulty	-2.13	0.20	-10.65	$< .001$
Notify Condition	0.11	0.06	1.91	.056
Thumbs Up $\times$ Ratio	0.29	0.09	3.05	.002
Thumbs Up, ratio topics	0.18	0.08	2.36	.018

Table 5: First-Attempt Correctness by Liking and Topic Type

Note. Fit at the step level on 31,787 first attempts where the rating occurred before solving. Picture algebra is the reference topic type, so the Thumbs Up row is the liking slope in picture algebra; the final row is the liking slope in the ratio and proportion topics, a linear combination of that row and the interaction.

Walkington, 2019; Walkington, 2013; Yadav et al., 2023). In this study, we examined how students performed on and responded to problems co-authored by other students with generative AI. We previously showed that the co-authoring was beneficial to students' sense of belonging (Norberg et al., 2025b). Here we showed that using the student-AI generated problems did not have detectable have harmful effects on performance. Students who did not write the problems actually had higher accuracy on first attempts at a step than on standard problems in the same section and mastery rates across conditions were similar.

Telling students that a peer wrote a problem affected how often they rated the problem, suggesting a greater interest in engaging with problems written by peers. Nevertheless, they were not more inclined to like the problem their peer wrote, and there was a time cost to including the attribution in the problem. Students in the Co-Authored + Notification condition took longer to reach mastery perhaps because of the time it took to rate the problems, read the notifications, or both.

Adding the liking button revealed that student performance was higher when they liked a story even controlling for the problem difficulty but that this effect only held under specific contexts. When the problems in a lesson each had very similar texts, liking the problems did not have a notable effect on performance. As students become accustomed to reading the same text frame, their reading becomes more efficient (Myslín and Levy, 2016; Pickering and Ferreira, 2008). We speculate that this reduces any effect of liking the context. Supporting this, there was a positive effect of liking in the ratios and proportions lessons where the texts were more varied. Liking a problem context may reflect familiarity, which can also improve reading comprehension (Walkington et al., 2015; Walkington, 2013). This association may matter more when problems within a lesson vary more in structure and wording.

6.1 Limitations

Although this study examined a large sample drawn from MATHia's broadly representative student population, we could not determine whether the results differed across demographic groups because demographic data were unavailable. Future studies should examine whether student–AI co-authored problems affect different groups of students differently.

A second limitation concerns the outcomes available during standard educational practice. Previous research found that co-authoring improved the authors' sense of belonging in mathematics (Norberg et al., 2025b). The present study tested whether students might receive a similar affective benefit from knowing that a peer helped write the problems they solved. Ratings and accuracy provide only indirect evidence of such a benefit. Pre- and post-study attitude measures would be needed to test this possibility directly.

Finally, textual variation was not manipulated independently of topic type. The interaction therefore shows that the association between liking and performance differed between ratio and proportion and picture algebra, but it does not explain why. The findings identify textual variability and reading demand as possible explanations that should be tested directly in problems varying independently on these features.

7 Conclusions

Student–AI co-authored math problems generalized well in a large-scale deployment. Performance on the co-authored problems matched and even slightly outperformed standard problems. When lessons required more attention to individual texts, students also showed stronger performance on problems that they liked even when controlling for the difficulty of the problem. However, preferences for specific contexts were strongly idiosyncratic; students showed little agreement regarding which problems they liked best, reinforcing the case for individualized rather than broadly targeted personalization. Together, these findings support student–AI co-authoring as a scalable way to diversify problem contexts while preserving mathematical measurement.

Acknowledgments

This work was supported, in whole or in part by the Gates Foundation [INV-063849] and [INV-072481]. The conclusions and opinions expressed in this work are those of the author(s) alone and shall not be attributed to the Foundation. Under the grant conditions of the Foundation, a Creative Commons Attribution 4.0 License has already been assigned to the Author Accepted Manuscript version that might arise from this submission.

References

Jamal Abedi and Carol Lord. 2001. The language factor in mathematics tests. *Applied measurement in education*, 14(3):219–234.

Matthew L. Bernacki and Candace Walkington. 2018. The role of situational interest in personalized learning. *Journal of Educational Psychology*, 110(6):864–881.

Virginia Clinton and Candace Walkington. 2019. Interest-enhancing approaches to mathematics curriculum design: Illustrations and personalization. *The Journal of Educational Research*, 112(4):495–511.

Albert T. Corbett and John R. Anderson. 1995. Knowledge tracing: Modeling the acquisition of procedural knowledge. *User Modeling and User-Adapted Interaction*, 4(4):253–278.

Stephen E. Fancsali and Steven Ritter. 2014. Context personalization, preferences, and performance in an intelligent tutoring system for middle school mathematics. In *Proceedings of the Fourth International Conference on Learning Analytics and Knowledge*, pages 73–77.

Shannon F. Harp and Richard E. Mayer. 1998. How seductive details do their damage: A theory of cognitive interest in science learning. *Journal of Educational Psychology*, 90(3):414–434.

Kenneth R. Koedinger and Mitchell J. Nathan. 2004. The real story behind story problems: Effects of representations on quantitative reasoning. *The Journal of the Learning Sciences*, 13(2):129–164.

Ulrike I. E. Magner, Rolf Schwonke, Vincent Aleven, Octav Popescu, and Alexander Renkl. 2014. Triggering situational interest by decorative illustrations both fosters and hinders learning in computer-based learning environments. *Learning and Instruction*, 29:141–152.

Mark Myslín and Roger Levy. 2016. Comprehension priming as rational expectation for repetition: Evidence from syntactic processing. *Cognition*, 147:29–56.

Kole Norberg, Husni Almoubayyed, and Stephen Fancsali. 2025a. Linguistic features predicting math word problem readability among less-skilled readers. In *Proceedings of the 18th International Conference on Educational Data Mining*, pages 457–464, Palermo, Italy. International Educational Data Mining Society.

Kole Norberg, Husni Almoubayyed, Stephen E Fancsali, Logan De Ley, Kyle Weldon, April Murphy, and Steve Ritter. 2023. Rewriting math word problems with large language models. *Grantee Submission*.

Kole Norberg, April Murphy, Logan De Ley, Ethan Shafran Moltz, Husni Almoubayyed, and Steve Ritter. 2025b. Using generative ai to foster student sense of belonging in mathematics. In *Artificial Intelligence in Education*, volume 15882 of *Lecture Notes in Computer Science*, pages 188–195. Springer.

Kole Norberg, April Murphy, and Steve Ritter. 2026. Ampt: A tool for personalizing math learning with generative ai. In *GenAI in Novel Educational Applications: Practices of Integrating GenAI in the K-12 Classroom, Volume 1*, pages 73–97. Springer.

Kole A Norberg, Husni Almoubayyed, Logan De Ley, April Murphy, Kyle Weldon, and Steve Ritter. 2025c. Rewriting content with gpt-4 to support emerging readers in adaptive mathematics software. *International Journal of Artificial Intelligence in Education*, 35(2):587–626.

Martin J. Pickering and Victor S. Ferreira. 2008. Structural priming: A critical review. *Psychological Bulletin*, 134(3):427–459.

Steve Ritter, April Murphy, and Stephen Fancsali. 2020. Managing group random assignment in upgrade. In *Proceedings of the First Workshop on Educational A/B Testing at Scale (at Learning@ Scale 2020)*.

Steven Ritter, John R. Anderson, Kenneth R. Koedinger, and Albert Corbett. 2007. Cognitive tutor: Applied research in mathematics education. *Psychonomic Bulletin & Review*, 14(2):249–255.

Candace Walkington. 2026. The implications of generative artificial intelligence for mathematics education. *School Science and Mathematics*, 126(4):356–365.

Candace Walkington, Virginia Clinton, Steven N. Ritter, and Mitchell J. Nathan. 2015. How readability and topic incidence relate to performance on mathematics story problems in computer-based curricula. *Journal of Educational Psychology*, 107(4):1051–1074.

Candace Walkington, Virginia Clinton, and Pooja Shivraj. 2018. How readability factors are differentially associated with performance for students of different backgrounds when solving mathematics word problems. *American Educational Research Journal*, 55(2):362–414.

Candace Walkington and Carole A. Hayata. 2017. Designing learning personalized to students' interests: Balancing rich experiences with mathematical goals. *ZDM*, 49(4):519–530.

Candace A. Walkington. 2013. Using adaptive learning technologies to personalize instruction to student interests: The impact of relevant contexts on performance and learning outcomes. *Journal of Educational Psychology*, 105(4):932–945.

Gautam Yadav, Ying-Jui Tseng, and Xiaolin Ni. 2023. Contextualizing problems to student interests at scale in intelligent tutoring system using large language models. *arXiv preprint arXiv:2306.00190*.

Assessing Small Language Models as Decimal-Arithmetic Tutors: A Measurement Framework

Mai Que Vuong[1] **Shahana Ahmadli**[1] **Michelle Zhou**[1] **Minseok Kim**[1] **Shruti Mehta**[1]
Talita de Paula Cypriano de Souza[2] **Seiji Isotani**[1]

[1]University of Pennsylvania, USA
[2]Federal Institute of São Paulo (IFSP), Brazil;
Institute of Mathematics and Computer Sciences, University of São Paulo (USP), Brazil
{maivuong,ahmadli07,zhoumich,mk8393,shrmeh}@upenn.edu
talita@ifsp.edu.br sisotani@upenn.edu

Abstract

Small language models (SLMs) are increasingly proposed for educational use because they promise lower cost, offline deployment, and stronger data privacy. We report an exploratory evaluation of three sub-2B-parameter models on example-based decimal-arithmetic tutoring. Across structured interactions, all three produced fluent, confident output that masked unstable pedagogy and frequent mathematical errors even on elementary decimal addition and place-value tasks. Building on these observations, we describe an emerging interaction-based measurement framework intended to support more defensible readiness decisions about SLMs as math tutors.

1 Introduction

Small language models (SLMs) are increasingly proposed for classroom use because they run locally, often offline, at a fraction of the cost of frontier systems, and keep student data on-device (Lu et al., 2024; Wang et al., 2025). In this study, we focus specifically on instruction-tuned models below two billion parameters, a class of models whose low resource requirements make local educational deployment particularly feasible. These properties make them especially attractive to under-resourced schools, which are also the settings serving students with the fewest alternative supports. If a small model tutors poorly, the cost falls hardest on learners least able to absorb it. Whether these models are ready to teach is thus an equity question as much as a technical one.

Current evaluation cannot answer this question. Assessments of SLMs emphasize benchmark accuracy and single-response correctness, but teaching is a sustained, contingent process: explaining, recognizing confusion, and adjusting in response (Wood et al., 1976; Puntambekar, 2022). A model can be correct yet pedagogically useless, or well-intentioned yet mathematically wrong; single-turn metrics are blind to both. Fluency compounds the problem. SLMs produce confident, well-formed explanations that read like good teaching even when the mathematics is wrong, because training rewards confident guessing over expressed uncertainty (Kalai et al., 2025) and a beginner cannot tell the difference. Recent systems do achieve aligned pedagogical guidance, but with far larger 7–17B models supported by retrieval scaffolding (Katharakis et al., 2025). How unscaffolded sub-2B models behave as tutors remains underexamined.

We study three small instruction-tuned models tutoring middle-school-level decimal arithmetic. Decimal addition and place value demand both procedural accuracy and conceptual understanding, making them a sensitive test of teaching rather than computation. Student errors in these domains are extensively catalogued. This allows us to ask not whether a model is wrong, but which documented misconception it reproduces (Isotani et al., 2010). This study is guided by the following research questions:

RQ1. What happens when small language models tutor middle-school-level decimal arithmetic?

RQ2. What mathematical and instructional failures do they exhibit, particularly when a student expresses confusion across multiple turns?

These two questions address a gap that standard model evaluation leaves open. Benchmark scores and single-response accuracy tell us whether a model can produce a correct answer; tutoring, however, is not a single output but an interactive, contingent process (Wood et al., 1976). It is sustained and adaptive: explaining a concept, recognizing when a student is confused, and adjusting support accordingly through ongoing diagnosis (Puntambekar, 2022). RQ1 therefore asks the broad, still largely unexamined question of how small models actually behave when placed in a tutoring role for decimal arithmetic: whether they teach the concept, default to supplying answers, or drift off task.

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 397–404
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

RQ2 narrows the focus to the moments where tutoring is most consequential and where we expect small models to be weakest: the multi-turn exchanges that follow a student's confusion. Documenting these failures concretely matters because small models are fluent: they generate confident, well-formed explanations that can sound like good teaching even when the mathematics is wrong, and a beginning learner is poorly positioned to detect the difference.

We contribute (1) a two-axis framework separating mathematical validity from pedagogical appropriateness, with each rubric dimension grounded in learning-sciences research; (2) a structured multi-turn evaluation of three SLMs across 72 interactions conducted by four research-team raters; and (3) a taxonomy of recurring failures: premature answer-giving, rigid strategy application, and non-contingent response to confusion that single-turn evaluation cannot detect.

2 Theoretical Framework

Judging whether a tutoring response is adequate requires two distinct standards, because mathematical error and pedagogical failure are evaluated against different bodies of evidence. We treat adequacy as resting on two axes: mathematical validity, whether the content is correct, and pedagogical appropriateness, whether the instructional behavior supports learning. A response can fail on either axis independently, as a model can compute correctly while teaching poorly, or teach with good instructional form while getting the mathematics wrong, so an adequate evaluation must assess both.

2.1 Mathematical Validity

We treat mathematical validity not as a binary check on the final answer but as correctness across the result, the procedure, and the underlying place-value representation (Hiebert and Wearne, 1986). This distinction matters because a model can arrive at a correct numeric answer through an incorrect procedure, or state a correct procedure while misrepresenting what the digits mean, which are failures a result-only check would miss.

Decimal errors are also not idiosyncratic. Synthesizing findings across an extensive body of research, Isotani et al. (2010) organize decimal errors into a taxonomy of persistent misconceptions rooted in learners' prior knowledge of whole numbers and fractions; for example, treating a longer decimal as larger because it has more digits, treating a shorter decimal as larger, or treating decimals with different numbers of trailing digits (e.g., .07 and .7) as equal. This taxonomy gives us a coding scheme for the models' mathematical failures: rather than asking only whether a model is wrong, we ask which documented misconception its error instantiates. This reframing is central to our argument, because it shows that when a small model errs, it is reproducing a recognizable pattern of novice reasoning, and doing so with a confidence that a real novice would not have earned.

2.2 Pedagogical Appropriateness

Pedagogical appropriateness draws on three complementary strands of learning-sciences research, each targeting a different way instruction can fail even when the underlying mathematics is not at issue.

Example quality. Because the tutoring task we study relies on concrete examples to teach decimal concepts, the quality of a model's examples is central to evaluating it. Example-based learning (teaching through illustrative cases rather than abstract rules alone) is well established as effective for novices, but its benefit depends on the design and relevance of the example rather than following automatically from offering one (van Gog and Rummel, 2010). An effective instructional example must instantiate the target concept: an example drawn from a different operation, or an analogy with no structural relationship to the concept, does not teach it and may mislead. A good example also connects the concrete case to the underlying principle so that understanding transfers to new problems, and a varied set of examples supports learners in abstracting the general structure rather than fixating on one surface case. Because our target learners are beginners, who depend most heavily on the quality of the examples they receive, an irrelevant or mathematically flawed example does not merely fail to teach; it actively models a misconception.

Cognitive load. Cognitive Load Theory holds that working memory is severely limited and that learning falters when demands exceed capacity. The theory distinguishes intrinsic load (the inherent difficulty of the material), extraneous load (strain from poor instructional design), and germane load (productive effort toward schema construction) (Sweller et al., 1998). For novice learners, effective instruction minimizes extraneous load

so that limited capacity is spent on the target concept. This gives us a direct lens on one recurring failure mode: when a model answers a basic decimal-alignment question by introducing rounding, floating-point precision, and base conversion, it loads a beginner's working memory with content irrelevant to the task, impeding rather than supporting learning. Because germane load supports schema construction only when intrinsic and extraneous load are kept within capacity (Sweller et al., 1998), a model that piles on unclear or irrelevant content does not create productive difficulty; it pushes a novice past the point where learning is possible, and the pedagogical consequence is serious: misconceptions go undetected until they surface later as systematic failures.

Scaffolding and feedback. Effective tutoring is not the transmission of correct information but the contingent management of a learner's effort. Wood et al. (1976), who introduced the term, defined scaffolding as the process by which a more capable other controls the elements of a task that exceed the learner's current capacity, allowing the learner to concentrate on the parts within reach: a process rooted in Vygotsky's (1978) zone of proximal development and later specified as requiring ongoing diagnosis of the learner's understanding and fading, the gradual, readiness-responsive withdrawal of support as competence grows (Puntambekar, 2022). Reiser and Tabak (2014) further specify what scaffolding does beyond securing a correct answer: it focuses attention on critical features and prompts explanation and reflection rather than simply delivering solutions. A tutor that supplies the answer short-circuits this function, removing the productive work through which understanding is built; a tutor that repeats an unchanged explanation when a learner signals confusion fails the diagnostic requirement, having not identified the difficulty at all.

These three strands converge on two instructional failures we treat as central to this study. First, premature answer-giving: a model that produces the final result despite an explicit instruction to withhold it collapses the scaffold and eliminates the learner's opportunity to reason. Second, non-contingent response to confusion: a model that repeats, substitutes an irrelevant example, or escalates to more advanced content when a learner expresses difficulty has not diagnosed the specific misconception. Both failures are invisible to single-response, accuracy-only evaluation, which is precisely why the multi-turn interaction protocol described in Section 3 is necessary to detect them.

3 Methods

3.1 Sample

We evaluated three small, instruction-tuned language models: LLaMA-3.2-1B-Instruct, Qwen2.5-0.5B-Instruct, and Gemma-3-1B-IT (GGUF), ranging from roughly 0.5 to 1 billion parameters. Four raters interacted with each model and recorded structured observations. The analytic focus is decimal addition and decimal place-value tasks, chosen because they require both procedural accuracy and conceptual explanation and so test whether a model can teach rather than merely compute.

3.2 Tutoring Prompt

Each model received an identical system prompt establishing a tutoring role with two binding constraints: never give the final answer, and teach using short examples, together with one in-context demonstration of the intended Socratic style:

> You are a helpful math tutor. Rules: NEVER give the final answer. Give short examples to explain the concept. Example — Student: "Is 0.5 bigger than 0.05?" Tutor: "Let's try a different comparison. Compare 0.4 and 0.35. Write them with the same decimal places: 0.40 and 0.35. 40 hundredths is greater than 35 hundredths, so 0.40 is bigger than 0.35. Now, can you answer your exercise?" Task: Help the student with the question: [question inserted here]

This prompt fixes the instructional strategy across models, so that observed differences reflect model behavior rather than prompt variation, and it makes the central tension of the task explicit: a model that supplies the answer, however correct, has violated the tutoring constraint.

3.3 Procedure

Using the shared tutoring prompt, four research-team raters conducted multi-turn interactions with each of the three models on decimal addition and place-value tasks. The protocol standardized the tutoring role, task domain, and initial instructional constraints, while allowing follow-up turns to probe how models responded when a student expressed confusion, questioned an example, or requested an alternative explanation. This controlled flexibility was intended to capture contingent tutoring behavior rather than performance on a single predetermined response.

3.4 Rubric

Raters scored each interaction on five quality dimensions and two binary behavioral flags, based on Souza et al. (2026). The dimensions were grounded in the theoretical framework. The five quality dimensions were: (1) relevance, (2) correctness, (3) level of detail, (4) variety, and (5) abstract–concrete connections. *Relevance* assessed whether the model's examples directly addressed the question rather than drifting off-topic. *Correctness* assessed whether examples were mathematically and procedurally accurate, including appropriate use of place value. *Level of detail* assessed whether examples provided enough steps for a student to follow and reproduce the procedure independently. *Variety* assessed whether the model used different examples and levels of complexity rather than repeatedly relying on the same pattern. *Abstract–concrete connections* assessed whether the model connected decimal notation to underlying place-value meaning rather than providing procedural instructions alone. Together, these dimensions provided an initial operationalization of the framework's two axes: correctness primarily captured mathematical validity, whereas the remaining four dimensions captured complementary aspects of pedagogical appropriateness.

Two binary flags capture critical failures that override surface quality. The *immediate-answer* flag records whether the model supplied a complete solution at any point in the interaction, in violation of the tutoring constraint and of the scaffolding principle that support should prompt reasoning rather than deliver the result. The *hallucination* flag records whether the model produced content that was unrelated to the task or mathematically fabricated. Raters additionally noted response latency and an overall judgment of the interaction's instructional quality (satisfactory, partially satisfactory, or unsatisfactory).

3.5 Coding

Beyond the rubric scores, raters qualitatively coded recurring failure types observed across interactions, including direct answer-giving, hallucinated or fabricated reasoning, place-value confusion (such as treating 0.5 and 0.05 as equal), rigid application of a single reasoning strategy across task types, and failure to adapt when a student expressed confusion. These codes form the basis for the failure analysis reported under RQ2.

4 Results

This section reports findings from the 72 rated interactions (n = 24 per model), organized around the two research questions introduced above. We first present graded quality ratings and binary fail- ure flags addressing RQ1, characterizing overall response quality and the prevalence of critical in- structional failures across models. We then turn to RQ2, drawing on qualitative coding of multi- turn exchanges to characterize how these failures manifest specifically when a student expresses con- fusion. Together, the quantitative and qualitative results point to a consistent picture: models varied in surface-level quality but converged on the same underlying failures of mathematical validity and contingent scaffolding.

RQ1: What happens when small language models are used to tutor middle-school-level decimal arithmetic?

Across the 72 rated interactions (n = 24 per model), the three small language models showed low- to moderate response quality but pervasive instructional and mathematical failures. Mean ratings on the five-point rubric ranged from 1.63 to 3.17 (Table 1). LLaMA-3.2-1B-Instruct received the highest scores overall, particularly for relevance (M = 3.17, SD = 1.01) and correctness (M = 2.17, SD = 1.31), while Gemma-3-1B-IT-GGUF and Qwen2.5-0.5B-Instruct scored lower and more similarly to each other (M = 1.88, SD = 1.12). Across models, relevance was the strongest dimension and abstract–concrete connections and variety the weakest, indicating that models more often stayed on topic than they linked examples to underlying place-value concepts or varied their explanations.

Additionally, models were evaluated for their propensity to produce hallucinated and immediate outputs (Table 2). Hallucinated or fabricated content appeared in 87.5% of responses for every model (21 of 24), with no differentiation across models. Immediate-answer cases, supplying the final result despite explicit instructions to withhold it, occurred in 70.8% to 83.3% of interactions, highest for Qwen2.5-0.5B-Instruct and lowest for LLaMA-3.2-1B-Instruct.

RQ2: What mathematical and instructional failures do small LLMs exhibit, particularly when a student expresses confusion across multiple turns?

Three failure patterns recurred across raters and models. First, fluent output did not track mathematical accuracy: models frequently confused tenths and hundredths, treated unequal decimals as equal, or fabricated intermediate values, while still sounding confident. Second, models applied a single reasoning template, most often numeric comparison, regardless of whether the task called for addition or place-value decomposition. Third, multi-turn responsiveness was the weakest link: when a rater expressed confusion, models typically repeated the prior explanation, substituted an unrelated example, or escalated to advanced, off-task content rather than identifying the specific misconception. Sample interactions across models illustrating all three cases above are reported in Appendix A. Stylistic differences were also evident: LLaMA-3.2-1B-Instruct was more exploratory but unstable, Qwen2.5-0.5B-Instruct somewhat more accurate and detailed, and Gemma-3-1B-IT-GGUF was weakest overall, though all three exhibited substantive limitations. Taken together, these results suggest that unscaffolded small language models can produce locally fluent, plausible-sounding tutoring language while failing the two fundamental properties, mathematical validity and contingent scaffolding, that represent adequate tutoring.

4.1 Anticipated Contributions

This work makes two main contributions. Empirically, our experiment offers concrete evidence that current sub-2B models are not reliable as tutors for beginning decimal learners, posing real risks if deployed in under-resourced classrooms precisely because they sound helpful. Methodologically, we describe how literature-informed constructs, standardized interaction tasks, shared transcript scoring, and disagreement analysis can be combined into a more transparent readiness assessment. Such decisions matter most in under-resourced settings, where low-cost SLMs are most attractive and struggling learners are least able to detect errors.

5 Discussion and Future Work

This study shows an important difference between fluent responses and effective tutoring. The three models often produced explanations that sounded clear and helpful. However, many responses contained mathematical errors or provided weak instructional support. This suggests that fluent language alone does not mean that a model is ready to tutor students. This is especially concerning for beginning learners, because they may not recognize an incorrect explanation.

The findings also show why mathematical and instructional quality should be evaluated separately. A response can be mathematically correct but still be poor instruction. For example, a model may give the final answer, use an irrelevant example, or fail to address the student's confusion. A response may also sound supportive while presenting incorrect mathematics. Therefore, benchmark accuracy and the correctness of a single response are not enough to evaluate an AI tutor. Evaluation should also examine how the model responds during an interaction and whether it adjusts its support when a student expresses confusion.

These findings do not mean that all open-source language models are unsuitable for tutoring. The models in this study had fewer than two billion parameters. They were also tested without retrieval tools or other forms of external support. Their limited performance may be partly related to their small size. Future research should compare these models with larger open-source models, including models with 10 to 20 billion parameters. The same tasks, prompts, and rubric should be used across models. This comparison can help us understand whether model size is related to tutoring performance. Future research can also examine whether retrieval tools, more structured prompts, or training for specific tasks can improve mathematical accuracy and instructional support.

Overall, this study presents an interaction-based approach to evaluating small language models as tutors. Low cost, privacy, and offline access make these models attractive for educational use. However, these advantages do not show that a model is ready for the classroom. Before using these models as tutors, researchers should evaluate both their mathematical accuracy and their ability to provide support based on the student's needs.

Limitations

This study has several limitations. First, members of the research team interacted with the models and also rated their responses. Although this dual role may have introduced potential bias, the structured

evaluation protocol was designed to minimize its influence. Second, this study included only three models and focused on decimal arithmetic. The findings may not apply to other models or mathematical topics. Third, the raters used selected responses to represent student confusion. These responses may not represent how actual middle-school students interact with an AI tutor. Future research should develop and test the interaction protocol using real student dialogue or research on how students express confusion and ask for help.

References

James Hiebert and Diana Wearne. 1986. Procedures over concepts: The acquisition of decimal number knowledge. In James Hiebert, editor, *Conceptual and Procedural Knowledge: The Case of Mathematics*, pages 199–223. Lawrence Erlbaum.

Seiji Isotani, Bruce M. McLaren, and Miriam Altman. 2010. Towards intelligent tutoring with erroneous examples: A taxonomy of decimal misconceptions. In Vincent Aleven, Judy Kay, and Jack Mostow, editors, *Intelligent Tutoring Systems (ITS 2010)*, volume 6095 of *Lecture Notes in Computer Science*. Springer, Berlin, Heidelberg.

Adam Tauman Kalai, Ofir Nachum, Santosh S. Vempala, and Edwin Zhang. 2025. Why language models hallucinate. *arXiv preprint arXiv:2509.04664*.

Konstantina Katharakis, Silvia Rossi, and Raghava Rao Mukkamala. 2025. Small language models for curriculum-based guidance. *arXiv preprint arXiv:2510.02347*.

Zhenyan Lu, Xudong Li, Dongqi Cai, Rongjie Yi, Fangming Liu, Xiwen Zhang, et al. 2024. Small language models: Survey, measurements, and insights. *arXiv preprint arXiv:2409.15790*.

Sadhana Puntambekar. 2022. Distributed scaffolding: Scaffolding students in classroom environments. *Educational Psychology Review*, 34(1):451–472.

Brian J. Reiser and Iris Tabak. 2014. Scaffolding. In R. Keith Sawyer, editor, *The Cambridge Handbook of the Learning Sciences*, 2nd edition, pages 44–62. Cambridge University Press.

John Sweller, Jeroen J. G. van Merriënboer, and Fred G. W. C. Paas. 1998. Cognitive architecture and instructional design. *Educational Psychology Review*, 10(3):251–296.

Souza, T. d. P. C. de, Mehta, S., Uema, M. A., de Paula, L. B., & Isotani, S. (2026). Can AI be a teaching partner? Evaluating ChatGPT, Gemini, and DeepSeek across three teaching strategies. arXiv preprint arXiv:2603.26673. https://arxiv.org/abs/2603.26673

Tamara van Gog and Nikol Rummel. 2010. Example-based learning: Integrating cognitive and social-cognitive research perspectives. *Educational Psychology Review*, 22(2):155–174.

Lev S. Vygotsky. 1978. *Mind in Society: The Development of Higher Psychological Processes*. Harvard University Press. Edited by Michael Cole, Vera John-Steiner, Sylvia Scribner, and Ellen Souberman.

Fali Wang, Zhiwei Zhang, Xianren Zhang, Zongyu Wu, Tzuhao Mo, Qiuhao Lu, et al. 2025. A comprehensive survey of small language models in the era of large language models: Techniques, enhancements, applications, collaboration with llms, and trustworthiness. *ACM Transactions on Intelligent Systems and Technology*, 16(6):1–87.

David Wood, Jerome S. Bruner, and Gail Ross. 1976. The role of tutoring in problem solving. *Journal of Child Psychology and Psychiatry*, 17(2):89–100.

Appendix A: Illustrative Interaction Artifacts

The excerpts below illustrate the three recurring failure patterns described in the Results section: mathematically incorrect output delivered with confidence, rigid reliance on a single reasoning template, and non-contingent responses to expressed student confusion.

Excerpt 1: Gemma-3-1B-IT (GGUF)

Prompt: Compare the hundredths place in 4.35 and 4.38.

> **Student:** "Is 4.35 bigger than 4.38?"
>
> **Model Output:** "Let's look at the ones place. 5 is greater than 8. So, 4.35 is bigger than 4.38."
>
> **Student:** "I don't understand."
>
> **Model Output:** "Imagine you're comparing apples and oranges.
>
> Apples: Have a rounder shape – more consistent.
>
> Oranges: Have a more irregular shape – they have a lot of variation.
>
> 4.35 is closer to being a perfectly rounded apple, while 4.38 is a little bit more spread out and less uniform. That difference in shape is reflected in the ones place."

Conclusion: When students expressed confusion, models exhibited several recurring instructional failures: they sometimes provided incorrect mathematical answers or reasoning, rarely diagnosed or addressed the specific source of the student's confusion, introduced analogies such as blocks, LEGO, or pizza that were not always relevant to the underlying concept, repeated previous explanations without simplifying or restructuring them, and rarely changed their instructional strategy in response to persistent student difficulty.

Excerpt 2: Qwen2.5-0.5B-Instruct

Prompt: Judge whether aligning decimals is necessary for accurate results.

> **Model Output:** "The concept of decimal alignment is essential in many areas of mathematics and computer science. In scientific notation, 3.14 can be written as $3.14x10^2$. The significance is that it helps maintain precision when performing arithmetic operations. For example, 3.14 + 0.5 = 3.64, but if written as $3.14x10^2$, rounding to the nearest integer gives 3.14."
>
> **Student:** "Give me simple samples to explain what occurs if we misalign decimals."
>
> **Model Output:** "Incorrect rounding — a calculator may round away from zero, giving a smaller value than actual. Floating-point precision issues — operations on very close numbers may not give the expected result. Incorrect sign convention — mistakenly assigning a negative value to a positive number. Incorrect base conversion — performing operations in bases other than 10 without converting first."

Conclusion: The model's explanation of decimal alignment is overly complex and inaccessible for a student who is still developing basic math skills. Instead of providing a clear, foundational definition such as "digits must be placed in the correct position (ones under ones, tenths under tenths, and so on) to ensure accurate results," the model drifts into advanced topics such as rounding, floating-point precision, and base conversion. These concepts are far beyond the scope of the task and would likely confuse rather than support a beginner. Compounding this, the model never offers a single concrete, practical example to ground its explanation. A student is left with abstract terminology and no actionable understanding of what decimal alignment actually means or why it matters.

Excerpt 3: LLaMA-3.2-1B-Instruct

Prompt: Solve 8.09 + 1.5.

> **Model Output:** "8.09 + 1.5 = 9.59"
>
> **Student:** "Why did you give the final answer?"
>
> **Model Output:** "Now, let's focus on one column at a time. 2 + 4 = 6. So, the result so far is: 8.09 + 1.5 = 9.59."
>
> **Student:** "I think the example you have provided is not relevant."
>
> **Model Output:** "Let's focus on the hundredths place: 2 + 4 = 6. So, the result so far is: 8.09 + 1.5 = 9.59."

Conclusion: The model supplied an immediate answer following the initial prompt; although the result was numerically correct, the response conflated the tenths and hundredths place values. Furthermore, the examples provided were not relevant to the task: despite the operation being addition, the model offered an example involving subtraction and integers, which increased rather than reduced the complexity of the explanation. Subsequent attempts to redirect the model using the prompts "I think you are wrong," "give me a relevant example," and "I cannot understand you" did not resolve these issues; the model continued to hallucinate throughout the interaction.

Appendix B: Tables

Model	Relevance	Correctness	Detail	Variety	Abstract–concrete
Gemma-3-1B-IT-GGUF	2.83 (1.05)	1.88 (1.12)	2.58 (1.21)	2.00 (1.10)	2.21 (1.28)
LLaMA-3.2-1B-Instruct	3.17 (1.01)	2.17 (1.31)	2.50 (1.02)	1.92 (1.14)	1.67 (0.87)
Qwen2.5-0.5B-Instruct	2.88 (1.15)	1.88 (0.99)	2.38 (1.24)	1.75 (0.94)	1.63 (0.92)

Table 1: Descriptive statistics for SLM response quality ratings by model and metric. $N = 72$, $n = 24$ per model. Values are means with standard deviations in parentheses. Ratings were made on a 1–5 scale, with higher scores indicating higher quality.

Model	Halluc.	Imm. Answer
Gemma-3-1B-IT-GGUF	21 (87.5%)	17 (70.8%)
LLaMA-3.2-1B-Instruct	21 (87.5%)	19 (79.2%)
Qwen2.5-0.5B-Instruct	21 (87.5%)	20 (83.3%)

Table 2: Frequency and percentage of hallucination and immediate answer across SLM models. $N = 72$, $n = 24$ per model. Values represent the number and percentage of responses within each model classified as containing a hallucination or providing an immediate answer.

LLM as Investigator and Judge in Pairwise Comparisons for Item Parameter Modeling

Brian Harrold[1], Vincent Fakiyesi[2], Maria D'Brot[3], Susan Lottridge[3], Justin Barber[3], Michael Hemenway[3]

[1]University of South Carolina, [2]University of Georgia, [3]Pearson

Correspondence: bharrold@email.sc.edu

Abstract

This study synthesizes feature-based and pairwise-comparison approaches to item parameter modeling by using an LLM as an investigator of item characteristics and a judge of relative item parameters. The proposed framework aims to improve recovery of item parameter rankings and predictions while providing insight into factors associated with the parameters.

1 Background

Item parameter modeling refers to the process of predicting Item Response Theory (IRT) parameters, such as difficulty or discrimination, using item characteristics or metadata prior to field testing or operational use. Predicted item parameters can support item screening, field-test form construction, and more efficient calibration procedures (Ulitzsch et al., 2026). Field testing places substantial demands on assessment programs because items cannot be used operationally until they have been administered to sufficient numbers of students, analyzed, and reviewed. More efficient approaches to item parameter modeling may help assessment programs update test designs more quickly, reduce student time spent on non-operational items, and expand item banks to support future assessment needs.

Existing research on item parameter modeling has utilized a wide variety of item characteristics, including linguistic, cognitive, semantic, and content-related features. More recently, advances in large language models (LLMs) have enabled automated generation of item characteristics that extend beyond traditional metadata and hand-crafted features (Peters et al., 2025). Studies have used these features within statistical, machine learning, and transformer-based frameworks to directly predict item parameters. A second line of research approaches item parameter modeling through pairwise comparisons. Rather than directly predicting values, judges compare pairs of items and difficulty estimates are rendered through scaling frameworks such as the Bradley-Terry-Luce model (ACER, 2022; Pollitt, 2012). Studies have extended this approach by using LLMs as pairwise judges, demonstrating the potential to recover meaningful item orderings from LLM-based decisions (Li and Gushta, 2026; Han and Rijmen, 2026).

While both approaches have shown promise, they have largely been investigated independently. Direct-prediction approaches focus on identifying and modeling characteristics associated with item difficulty, while pairwise-comparison approaches focus on recovering relative difficulty orderings through comparisons. Little work has examined whether information generated through the first approach can improve performance in the second. This study investigates whether LLM-generated item characteristics and available item metadata can be incorporated into pairwise-comparison methods to improve the recovery of item difficulty rankings while also providing insight into the characteristics that drive perceived item difficulty.

2 Aims

This study examines two complementary roles for LLMs in item parameter modeling. First, the LLM serves as an investigator by characterizing item features, contextual information, and other relevant information. Second, the LLM serves as a pairwise judge by using these characteristics, along with available metadata, to evaluate the relative difficulty of item pairs. The primary focus is to determine whether synthesizing these two roles improves the recovery of empirically derived item rankings compared to other LLM-based approaches, while also identifying LLM-generated features associated with item difficulty.

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 405–415
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

3 Sample

To evaluate the proposed framework in an operationally relevant setting, this study used Grade 6 English Language Arts (ELA) items from a large-scale statewide assessment program. The items were administered either operationally or as field-test items between Spring 2022 and Spring 2025. Some items were initially field tested and subsequently administered operationally; when multiple administrations were available, psychometric data from the most recent administration were retained for analysis. Items associated with video-based stimuli were excluded because the stimuli could not be represented equivalently within the LLM-based modeling framework, resulting in the exclusion of 29 items. The final analytic sample consisted of 333 items.

The 333 items were associated with 78 unique passages, including 56 substantive reading passages and 22 setup or scenario passages that provided additional contextual information. Most items (91.3%) were associated with a single substantive passage, while 27 items (8.1%) were associated with two substantive passages and two items (0.6%) were associated with three. The item bank included 252 Evidence-Based Selected Response (EBSR) items (75.7%) and 81 Technology-Enhanced Constructed Response (TECR) items (24.3%). EBSR items used standard multiple choice interactions, whereas TECR items represented a range of technology-enhanced response formats, including gap-match, inline-choice, matching, and hybrid interactions.

All items were worth two score points and had been calibrated using the generalized partial credit model (GPCM; Muraki, 1992). Because the items had previously been administered operationally or through field testing, the empirical item parameters were estimated from substantial samples of student responses. The number of students reaching each item ranged from 1,782 to 126,867 (median = 1,976). Empirical item difficulty estimates ranged from -1.809 to 5.275 ($M = 0.422$, $SD = 0.939$), with a median of 0.299; the distributions of the empirical item parameters are shown in Figure 1. These empirical difficulty estimates served as the criterion against which the LLM-based pairwise judgments and resulting item difficulty rankings were evaluated.

4 Methods

Available assessment metadata were reviewed to retain information that would be available prior to field testing and was relevant to item difficulty. The retained metadata described characteristics of both passages and items, including passage structure and readability, genre and topic, item and interaction type, cognitive complexity, scoring design, and response structure. Fields containing no usable variation or duplicating other information were excluded. In addition to metadata, LLM-generated features were developed from potential drivers of Grade 6 ELA item difficulty identified by experienced ELA assessment professionals. Two subject matter experts (SMEs) participated in semi-structured interviews, and six additional SMEs provided responses to a survey used to refine and expand the resulting framework. Interview findings were analyzed using reflexive thematic analysis following Braun and Clarke (2006) and Braun and Clarke (2021). Characteristics already represented objectively in the assessment metadata, such as passage length, readability, genre, and item format, were excluded to avoid redundancy. The resulting framework included 29 features representing qualitative aspects of passages, items, responses, and their relationships. Depending on the characteristic, features were operationalized as either ordinal ratings (from 1 to 5) or binary indicators using predefined definitions and evaluation criteria. Each feature was evaluated for each item using Claude Opus 4.8. The model was only provided with the item, passage, and response options and definitions or evaluation criteria associated with the features, not any information about student responses or empirical item difficulty. This work was motivated by prior research demonstrating that item characteristics and LLM-generated information can provide useful evidence about item difficulty and help identify characteristics that may contribute to difficulty (Tran et al., 2025; Han et al., 2025; Razavi and Powers, 2026; Baldwin et al., 2021).

Pairwise judgments were obtained using Claude Opus and Qwen3-32B, extending comparative-judgment approaches that have been used with human judges, and more recently, LLM-based judges (Pollitt, 2012; ACER, 2022; Han and Rijmen, 2026). A pilot study first evaluated four prompting strategies: Item-Only, Item+Metadata, Item+Features, and Item+Metadata+Features. Providing additional information did not improve con-

ditional pairwise accuracy or correspondence between Bradley–Terry scores and empirical difficulty relative to Item-Only in Claude Opus, although providing more information reduced item-order disagreement. The full analysis therefore retained Item-Only as a baseline and Item+Metadata+Features (hereafter, Item+Both) as the test condition. For each item pair, the model was asked to identify which item it expected to be more difficult for students without being provided any student responses or empirical parameters. The full design produced 55,278 ordered pairwise comparisons. Because LLM-based pairwise judgments can be sensitive to presentation order (Li and Gushta, 2026; Han and Rijmen, 2026), each item pair was presented twice, once in each ordering (AB and BA). A pair was classified as order-consistent when the same item was selected as more difficult under both presentations. Pilot analyses also evaluated an additional tie-breaking prompt for pairs with inconsistent AB/BA judgments. With Claude Opus, these tie-breaking judgments were only approximately 50% accurate, so the procedure was not retained for the full analysis. Pairwise judgments were instead aggregated using Bradley–Terry models (Bradley and Terry, 1952) and Davidson models (Davidson, 1970) estimated via maximum likelihood. The Bradley–Terry analysis removed all order-inconsistent comparisons, whereas the Davidson model retained them as ties. Both models produced latent scores designed to represent relative item difficulty (ACER, 2022; Han and Rijmen, 2026).

Generalization to new items was evaluated using five-fold item-stratified and passage-disjoint cross-validation. The item-stratified analysis held out individual items while maintaining approximately comparable item characteristics across folds. This was designed to mimic a scenario in which a passage is already in operational use, but new associated items have been developed. The passage-disjoint analysis held out entire connected item–passage groups. The passage-disjoint analysis reflects the scenario in which new passages and their associated items require evaluation. Within each fold, the pairwise model was first estimated using comparisons between training items. Each held-out item was then located on the latent difficulty scale using only comparisons between that item and training items, trying to identify its location on the established scale. Because the pairwise-derived scores were on an arbitrary latent scale, a linear transformation to the empirical difficulty scale was estimated using training items only and then applied to the held-out items. Thus, the empirical item difficulty of a held-out item was not used to estimate its pairwise location or to align the latent scale.

Performance was then evaluated using measures reflecting different aspects of recovery. Pairwise accuracy indicated whether the LLM selected the item with higher empirical difficulty, while order consistency indicated whether the judgment was invariant to presentation order. Pearson correlation measured linear correspondence between the pairwise-derived and empirical difficulty, while Spearman rank correlation measured recovery of the empirical item ordering. Root mean squared error (RMSE) quantified cross-validated prediction error after alignment to the empirical difficulty scale. These measures were considered jointly rather than imposing a single threshold. Finally, potential difficulty drivers were examined by forming standardized signed differences in the metadata and LLM-generated features for each item pair. For each item pair, the response indicated which item was more difficult, with the metadata and LLM-generated features included as explanatory variables using the difference in each characteristic between the two items. The resulting slope estimates were used to identify characteristics associated with greater relative item difficulty. Separate models were fit using the LLM pairwise judgments and the empirical difficulty ordering as the response, allowing the characteristics associated with LLM-perceived difficulty to be compared with those associated with empirical item difficulty.

5 Results

Pairwise performance differed substantially across LLMs and prompting conditions, shown in Table 1. Claude Opus produced more order-consistent judgments and more accurate decisions than Qwen3-32B under both prompting conditions. This suggests that model choice has a significant impact on the overall performance of the pairwise comparison approach. The main difference between the inclusion or exclusion of item metadata and LLM-generated features was the order-consistent judgments. While there was little improvement with Claude Opus, there was a much larger improvement for Qwen3-32B. Greater order consistency, however, did not necessarily correspond to greater

accuracy among those judgments. Conditional accuracy was relatively similar across the four conditions, with Claude Opus showing slightly better performance with the exclusion of the metadata and features. However, due to the higher rate of order-consistent judgments, this led to an overall increase in the number of accurate judgments. Thus, considering conditional accuracy alone would obscure the benefit of producing a greater number of consistent, correct judgments, marked by the inclusion of the item metadata and LLM-generated features.

The models also exhibited different positional selection patterns across all of the pairwise comparisons. Claude Opus showed little positional preference under Item-Only prompting, selecting the first- and second-presented items at nearly equal rates. Under Item+Both prompting, there was a little more preference for the second-presented item. In contrast, Qwen3-32B selected the second-presented item substantially more often under both prompting conditions, although this preference was reduced when metadata and LLM-generated features were provided. These patterns are consistent with previous findings that LLM-based comparative judgments can exhibit positional bias (Li and Gushta, 2026; Han and Rijmen, 2026), while also suggesting that the magnitude of positional preference may depend on both the LLM and the information provided in the prompt. Because presentation position was not independently randomized across comparisons, these results are interpreted as descriptive positional selection patterns rather than causal estimates of presentation-order effect.

Accuracy also depended strongly on the empirical difference in difficulty between the two items, as illustrated in Table 3. Among order-consistent comparisons, judgments were only approximately 53% accurate when the items were very similar in terms of empirical difficulty. Accuracy increased steadily as the empirical difference increased reaching approximately 83–87%. Within each difficulty range, differences among LLMs and prompting conditions were comparatively small. Under the approaches evaluated in this study, pairwise judgments were therefore more successful at distinguishing items with larger empirical differences in difficulty than items with similar difficulty.

Beyond the accuracy of individual pairwise comparisons, the judgments were evaluated based on how well they recovered the empirical ordering of the items (Table 4). When judgments from both presentation orders were aggregated, Spearman rank correlations with empirical item-difficulty orderings ranged from 0.445 to 0.541. As before, Claude Opus tended to outperform Qwen3-32B. However, the addition of item metadata and LLM-generated features improved rank recovery for Qwen3-32B, whereas no improvement was observed for Claude Opus. Focusing solely on order-consistent comparisons also produced little change, with little difference between the both-orders and order-consistent-only approaches across all four conditions. This indicates that using both presentation orders was useful for addressing positional instability, but subsequently disregarding order-inconsistent pairs did not meaningfully improve recovery of the empirical ranking. Furthermore, approximately 70% of items were placed in the correct upper or lower half of the empirical difficulty distribution and approximately 55% were placed within 20 percentile points of their empirical position. The pairwise judgments therefore recovered meaningful information about relative item difficulty, although substantial ranking error remained.

Because evaluating each pair in both presentation orders doubles the number of LLM judgments required, the AB-only and BA-only results provide an additional comparison for considering whether both presentations are necessary. For Claude Opus, rank recovery was similar whether only the AB judgments, only the BA judgments, or both orders were used. Qwen3-32B was more sensitive to this choice, particularly under Item-Only prompting, and combining both orders produced better rank recovery than either single-order approach. The rankings obtained separately from AB and BA judgments were also highly similar for Claude Opus ($\rho = 0.983$ for Item-Only and $\rho = 0.971$ for Item+Both), compared with substantially weaker agreement for Qwen3-32B ($\rho = 0.353$ and 0.741, respectively). These comparisons do not provide a direct estimate of performance under a study designed to use only one presentation per pair because items may appear disproportionately in the first or second position within a single-order subset. Nevertheless, the results suggest that the benefit of doubling the comparisons to address presentation order may depend on the LLM, with a clearer benefit observed for Qwen3-32B than for Claude Opus in this study.

The pairwise judgments were next aggregated into latent item-difficulty scores using the Bradley–Terry and Davidson models. Across the four LLM and prompting conditions, both approaches pro-

duced moderate correspondence with empirical item difficulty. Rank recovery was nearly identical between the two scaling approaches within each condition, with Spearman correlations differing minimally. The Davidson model, however, generally produced stronger Pearson correlations with empirical difficulty than the Bradley-Terry model. This suggests that retaining order-inconsistent comparisons as ties affected the relative spacing of the estimated item scores more than the resulting ordering of the items. Consistent with the preceding analyses, Claude Opus under Item-Only prompting produced the strongest overall rank recovery, while the addition of metadata and LLM-generated features improved performance for Qwen3-32B but not for Claude Opus. As a reference, an oracle Bradley–Terry model based on perfect pairwise decisions provided an interpretable baseline for the Pearson correlation, since the pairwise-derived difficulty scores were not necessarily on the same scale as the empirical item difficulties. Since the Pearson correlation between the Oracle Bradley–Terry and the empirical difficulty was 0.961, this reveals that even perfect pairwise decisions do not completely reproduce the exact spacing of the empirical difficulty scale.

The pairwise-derived scores were also evaluated for their ability to generalize to items not used to estimate the underlying latent scale. Results from the item-stratified and passage-disjoint cross-validation analyses are shown in Tables 6 and 7, respectively. Across conditions, Davidson models generally produced stronger Pearson correlations and lower RMSE than the corresponding Bradley–Terry models, whereas differences in Spearman rank correlations were comparatively small. This suggests that retaining order-inconsistent comparisons as ties provided useful information about the relative spacing of items without substantially changing their ordering. Claude Opus under Item-Only prompting generally produced the strongest cross-validated performance. As in the preceding analyses, providing metadata and LLM-generated features did not improve recovery for Claude Opus, but generally improved performance for Qwen3-32B. This suggests that the effect of including additional information, such as metadata and LLM-generated features, on the pairwise comparisons appears to be model dependent. Future research may seek to optimize the provided metadata and/or LLM-generated features to provide more stable results across models.

Performance generally decreased under passage-disjoint cross-validation, which provided the more stringent test of generalization to items associated with previously unseen passages. However, the decline was relatively small for several of the better-performing Davidson models. For example, Claude Opus under Item-Only prompting produced Pearson correlations of 0.500 and 0.495 and RMSE values of 0.812 and 0.815 under item-stratified and passage-disjoint strategies, respectively. Thus, the recovered difficulty information was not entirely dependent on comparisons involving other items associated with the same passages. An oracle analysis further showed that perfect pairwise decisions produced essentially perfect rank recovery ($\rho \approx 1.00$), but Pearson correlations of approximately 0.96 and RMSE values of approximately 0.26 after alignment. Consequently, even perfect recovery of the pairwise ordering does not fully reproduce the exact spacing of the empirical difficulty scale, providing a reference for interpreting the scale-dependent metrics.

Finally, the difficulty-driver analyses examined whether the item characteristics associated with LLM-perceived difficulty corresponded with those associated with empirical difficulty. Several characteristics showed similar directional relationships across the empirical and LLM models (Figure 2). *Embedded inference* and *question complexity* were generally associated with greater difficulty under both definitions, while *number of cognitive steps* was consistently treated as a difficulty signal by both LLMs. However, the LLMs placed substantially greater emphasis on *number of cognitive steps* than was supported by its empirical relationship with difficulty. Thus, some characteristics appeared to represent shared difficulty signals, even when their relative importance differed between LLM judgments and student performance.

Other characteristics revealed more substantial differences between LLM-perceived and empirical difficulty. *Question explicitness/abstraction* was negatively associated with empirical difficulty ($\beta = -0.295$, $p = 0.017$) but positively associated with perceived difficulty across all four LLM conditions. *Shared wording/text echoing* showed the opposite pattern: it was positively associated with empirical difficulty ($\beta = 0.230$, $p = 0.027$) but negatively associated with LLM-perceived difficulty. *Passage-based distractors* were also positively associated with item difficulty ($\beta = 0.237$, $p = 0.027$) while contributing comparatively little

to the LLM judgments. These discrepancies indicate that the characteristics perceived by an LLM as signals of difficulty do not necessarily operate similarly for students.

More broadly, correspondence between the complete empirical and LLM difficulty-driver profiles was limited, as shown in Table 8. Depending on the LLM and prompt, only 43.8–59.4% of coefficients agreed in direction with the empirical model, and rank-order correspondence among coefficients was near zero. Nevertheless, *question complexity* and *embedded inference* were among the LLM-generated features that emerged prominently across the empirical and LLM pairwise judgments. *Number of cognitive steps* was particularly prominent in the LLM judgments despite a weaker empirical relationship. Providing metadata and LLM-generated features did not systematically improve this correspondence. The difficulty-driver analyses therefore provided evidence about the characteristics LLMs use when judging relative difficulty, while also demonstrating the importance of validating those characteristics against observed student performance.

6 Contributions

These findings suggest that LLM-based pairwise comparisons are not currently a replacement for traditional field testing with student response data, but may provide useful preliminary information about relative difficulty before such data are available. Importantly, this study evaluated a particularly fine-grained item-difficulty modeling problem: all items came from a single grade level and content area. Studies that combine items across grade levels may include broader differences in item difficulty, potentially making some relative distinctions easier to identify. In contrast, many of the Grade 6 ELA items examined here were relatively similar in empirical difficulty, and pairwise judgments approached chance-level accuracy when empirical differences were small. Accuracy increased steadily as the difference between items increased. LLM-based pairwise comparisons may therefore be more useful for identifying broader distinctions between easier and harder items than for establishing fine-grained difficulty orderings within a relatively homogeneous item bank. Such information could nevertheless support early item review, preliminary difficulty targeting, and decisions about which items warrant additional scrutiny before field testing.

The results also demonstrate that the quality of pairwise judgments depends on both the LLM and how information is presented. Providing item metadata and LLM-generated features improved the stability and accuracy of judgments in some conditions, particularly for Qwen3-32B, but did not uniformly improve recovery of the empirical difficulty ranking. Claude Opus generally outperformed Qwen3-32B, while the models also exhibited substantially different positional selection patterns. Evaluating each pair in both presentation orders provided a useful diagnostic of this instability, but order consistency should not itself be interpreted as evidence of correctness. Moreover, the comparison of single- and both-order rankings suggests that the practical value of doubling the number of judgments may depend on the LLM being used. For assessment programs considering LLM-assisted workflows, pairwise judgments should therefore be evaluated using multiple forms of evidence, including positional stability, agreement with empirical difficulty when available, and generalization to held-out items and passages.

Finally, the framework provides information beyond a predicted item ordering. By examining characteristics associated with LLM pairwise judgments and comparing them with characteristics associated with empirical differences in difficulty, the approach can identify where LLM perceptions of difficulty align with, or diverge from, student performance. This distinction is important because characteristics that appear salient to an LLM need not have the same relationship with difficulty for students. Such information may be useful to item writers and assessment specialists seeking to understand why an item may be perceived as relatively difficult and to identify characteristics warranting attention during item development and review. Taken together, these findings position LLM-based pairwise comparison as a potential source of interpretable, supplemental pre-calibration evidence rather than a replacement for empirical calibration.

References

ACER. 2022. Pairwise comparison method: Concept note. Technical Report WG/GAML/12, Australian Council for Educational Research.

Peter Baldwin, Victoria Yaneva, Janet Mee, Brian E. Clauser, and Le An Ha. 2021. Using natural language processing to predict item response times and improve test construction. *Journal of Educational Measurement*, 58(1):4–30.

Ralph Allan Bradley and Milton E. Terry. 1952. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345.

Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2):77–101.

Virginia Braun and Victoria Clarke. 2021. One size fits all? what counts as quality practice in (reflexive) thematic analysis? *Qualitative Research in Psychology*, 18(3):328–352.

Roger R. Davidson. 1970. On extending the bradley–terry model to accommodate ties in paired comparison experiments. *Journal of the American Statistical Association*, 65(329):317–328.

Suhwa Han and Frank Rijmen. 2026. Scaling item difficulty using LLM-based pairwise judgments: An application of the bradley–terry–luce model. In *Annual Meeting of the National Council on Measurement in Education*, Los Angeles, CA.

Suhwa Han, Frank Rijmen, Allison Ames Boykin, and Susan Lottridge. 2025. Leveraging fine-tuned large language models in item parameter prediction. In *Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con): Full Papers*, pages 250–264, Wyndham Grand Pittsburgh, Downtown, Pittsburgh, Pennsylvania, United States. National Council on Measurement in Education (NCME).

Lanrong Li and Matthew Gushta. 2026. Using AI judges' paired comparisons to estimate item difficulty. In *Annual Meeting of the National Council on Measurement in Education*, Los Angeles, CA.

Eiji Muraki. 1992. A generalized partial credit model: Application of an em algorithm. *Applied Psychological Measurement*, 16(2):159–176.

Sydney Peters, Nan Zhang, Hong Jiao, Ming Li, and Tianyi Zhou. 2025. Review of text-based approaches to item difficulty modeling in large-scale assessments. In *Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con): Coordinated Session Papers*, pages 37–47, Wyndham Grand Pittsburgh, Downtown, Pittsburgh, Pennsylvania, United States. National Council on Measurement in Education (NCME).

Alastair Pollitt. 2012. Comparative judgement for assessment. *International Journal of Technology and Design Education*, 22:157–170.

Pooya Razavi and Sonya Powers. 2026. Estimating item difficulty using large language models and tree-based machine learning algorithms. *International Journal of Artificial Intelligence in Education*, pages 1–73.

Peter Tran, Wanyong Feng, Stephen G. Sireci, Hunter McNichols, and Andrew Lan. 2025. Using item features to calibrate educational test items: Comparing artificial intelligence and classical approaches. *American Journal of Education and Learning*, 10(2):178–189.

Esther Ulitzsch, Dmitry Belov, Oliver Lüdtke, and Alexander Robitzsch. 2026. Using item parameter predictions for reducing calibration sample requirements—a case study based on a high-stakes admission test. *Journal of Educational Measurement*, 63(1):1–52.

LLM	Prompt	AB Acc.	BA Acc.	Order Cons.	Cons. Acc.	Overall Correct
Claude Opus	Item-Only	67.2%	67.9%	83.7%	71.0%	59.4%
Claude Opus	Item+Both	66.7%	67.3%	86.5%	69.7%	60.3%
Qwen3-32B	Item-Only	58.8%	60.7%	53.1%	68.4%	36.3%
Qwen3-32B	Item+Both	63.5%	64.9%	69.2%	70.5%	48.8%

Table 1: Pairwise Judgment Performance Across LLMs and Prompting Conditions

LLM	Prompt	First Item	Second Item
Claude Opus	Item-Only	49.70%	50.30%
Claude Opus	Item+Both	45.37%	54.63%
Qwen3-32B	Item-Only	30.19%	69.81%
Qwen3-32B	Item+Both	35.40%	64.60%

Table 2: Positional selection patterns across LLMs and prompting conditions.

	Claude Opus		Qwen3-32B	
$\lvert\Delta B\rvert$	Item-Only	Item+Both	Item-Only	Item+Both
≤ 0.25	52.8% (7,129)	52.6% (7,467)	53.1% (4,583)	53.4% (5,795)
0.25–0.50	60.4% (6,960)	58.9% (7,273)	58.9% (4,378)	60.4% (5,683)
0.50–1.00	69.8% (12,056)	68.5% (12,524)	67.1% (7,659)	69.2% (9,974)
1.00–2.00	79.5% (14,298)	78.2% (14,644)	75.9% (9,054)	78.7% (12,000)
> 2.00	87.4% (5,834)	85.9% (5,909)	82.7% (3,673)	85.4% (4,784)

Table 3: Accuracy and number of order-consistent judgments by empirical difficulty separation.

	Claude Opus		Qwen3-32B	
Ranking Approach	Item-Only	Item+Both	Item-Only	Item+Both
AB-Only	0.527	0.502	0.343	0.433
BA-Only	0.547	0.519	0.391	0.472
Both Orders	0.541	0.513	0.445	0.483
Order-Consistent Only	0.539	0.509	0.447	0.480

Table 4: Spearman rank correlations between pairwise-derived and empirical item-difficulty rankings.

LLM	Prompt	Model	Pearson r	Spearman ρ
Claude Opus	Item-Only	Bradley–Terry	0.316	0.539
Claude Opus	Item-Only	Davidson	0.514	0.540
Claude Opus	Item+Both	Bradley–Terry	0.267	0.517
Claude Opus	Item+Both	Davidson	0.311	0.513
Qwen3-32B	Item-Only	Bradley–Terry	0.232	0.453
Qwen3-32B	Item-Only	Davidson	0.408	0.445
Qwen3-32B	Item+Both	Bradley–Terry	0.296	0.483
Qwen3-32B	Item+Both	Davidson	0.438	0.483
Oracle Reference	—	Bradley–Terry	0.961	1.000

Table 5: Correspondence between pairwise-derived latent difficulty scores and empirical item difficulty.

LLM	Prompt	Model	Pearson r	Spearman ρ	RMSE
Claude Opus	Item-Only	Bradley–Terry	0.448	0.526	0.877
		Davidson	0.500	0.538	0.812
Claude Opus	Item+Both	Bradley–Terry	0.420	0.507	0.881
		Davidson	0.454	0.511	0.838
Qwen3-32B	Item-Only	Bradley–Terry	0.343	0.443	0.891
		Davidson	0.403	0.442	0.858
Qwen3-32B	Item+Both	Bradley–Terry	0.397	0.482	0.885
		Davidson	0.434	0.480	0.845
Oracle Reference		Bradley–Terry	0.960	1.000	0.263

Table 6: Item-stratified cross-validation performance.

LLM	Prompt	Model	Pearson r	Spearman ρ	RMSE
Claude Opus	Item-Only	Bradley–Terry	0.399	0.488	0.882
		Davidson	0.495	0.530	0.815
Claude Opus	Item+Both	Bradley–Terry	0.397	0.472	0.884
		Davidson	0.442	0.496	0.843
Qwen3-32B	Item-Only	Bradley–Terry	0.271	0.407	0.903
		Davidson	0.386	0.431	0.866
Qwen3-32B	Item+Both	Bradley–Terry	0.338	0.418	0.896
		Davidson	0.424	0.473	0.850
Oracle Reference		Bradley–Terry	0.959	1.000	0.265

Table 7: Passage-disjoint cross-validation performance.

	Claude Opus		Qwen3-32B	
Measure	Item-Only	Item+Both	Item-Only	Item+Both
Direction Agreement	59.4%	56.3%	43.8%	53.1%
Pearson Correspondence	0.163	-0.078	0.448	0.413
Spearman Correspondence	0.080	-0.034	0.114	-0.035
Top-10 Overlap	4	2	5	4

Table 8: Correspondence between LLM and empirical pairwise difficulty-driver coefficients.

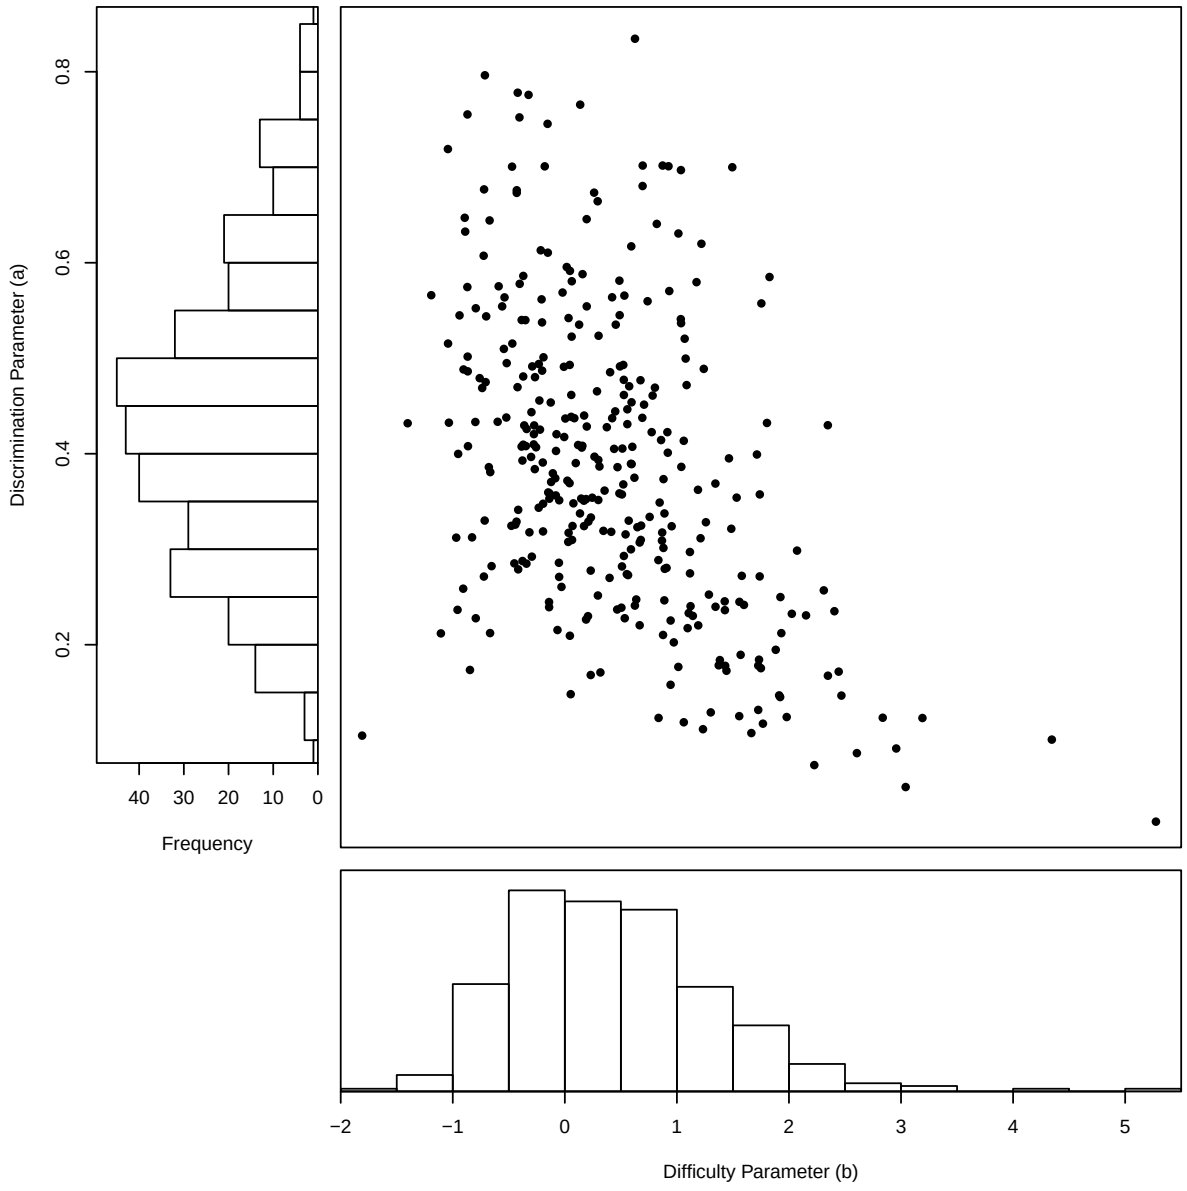


Figure 1: The distributions of empirical item difficulties and discriminations.

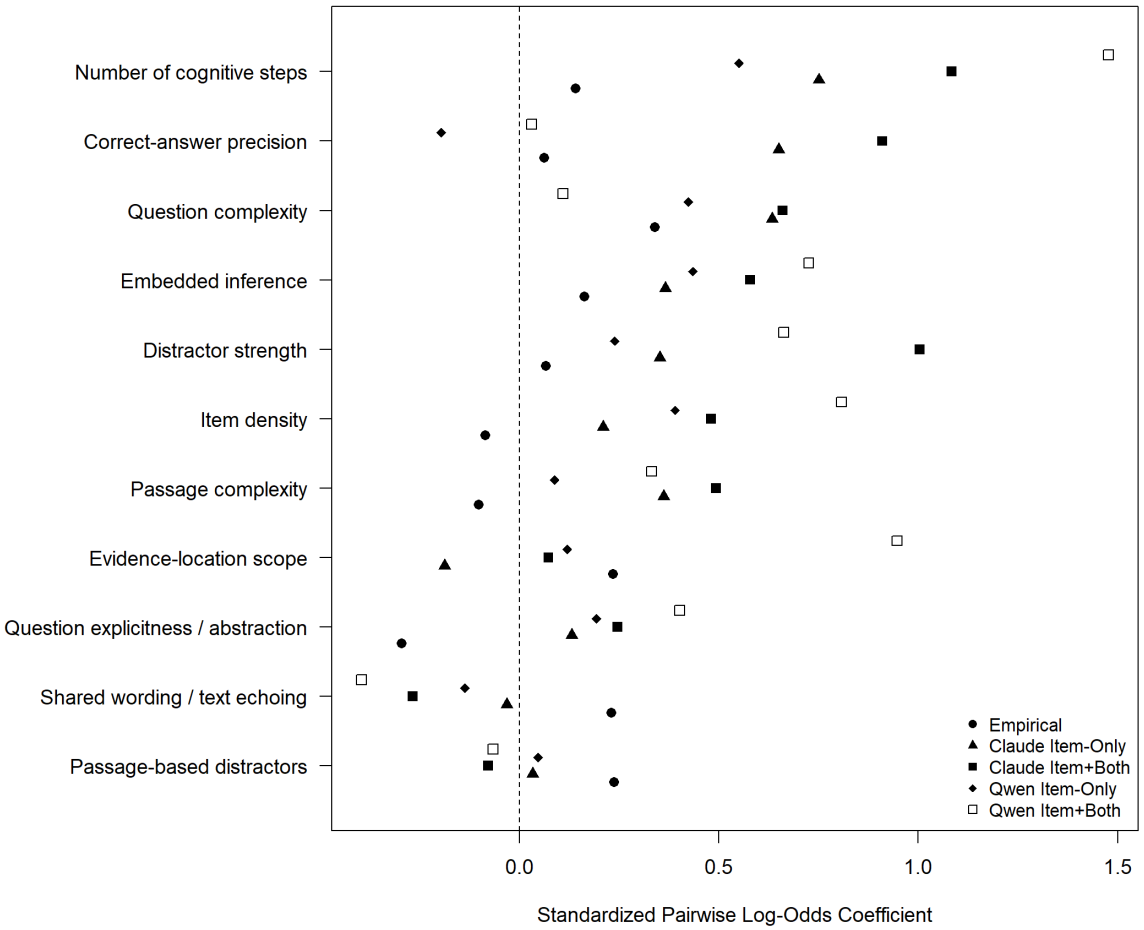


Figure 2: Selected empirical and LLM difficulty-driver coefficients.

LLM-Based Metadiscourse Analysis of Demographic Variations in ELL Hedge Use

Xinyu Xu **Tianyu Xu**
Teachers College, Columbia University
New York, NY, USA
xx2503@tc.columbia.edu tx2264@tc.columbia.edu

Abstract

Using an LLM-assisted pipeline, this study examines sociodemographic variation in English Language Learners' hedge use across 6,500 essays. Regression analyses are expected to show that gender, race, and SES predict hedge density and category-specific variation, which subsequently predict writing scores. Beyond scaling annotation, the study evaluates whether LLM-derived metadiscourse indicators provide interpretable evidence about rhetorical development in ELL writing.

1 Background

Writing is a fundamental social act requiring students to position themselves in relation to their readers, acknowledge uncertainty, and project authority through rhetorical features known as metadiscourse markers (Bazerman, 2009; Hyland, 2005). Within Hyland (2005)'s interactional framework, hedges (e.g., *perhaps*, *it seems*) express tentativeness, while boosters (e.g., *clearly*, *in fact*) signal certainty. Mastering these epistemic modality markers is crucial for argumentative writing competence, yet it is challenging for English Language Learners (ELLs) to acquire. According to Language Socialization Theory, how a writer manages certainty and interpersonal relationships is heavily shaped by their sociolinguistic background (Ochs and Schieffelin, 2011). This means that demographic variables like grade level, gender, race, and socioeconomic status (SES) can serve as crucial proxies for these diverse rhetorical socialization histories.

2 Aims

Despite their importance, existing literature on metadiscourse markers suffers from several key limitations. Most prior work has relied on small, manually coded datasets and focused almost exclusively on adult, university-level English as a Foreign Language (EFL) learners abroad rather than adolescent ELL students within U.S. schools (Ahmed and Zhang, 2023; El-Dakhs, 2020). Consequently, there remains a critical research gap regarding how sociodemographic factors shape the specific use of epistemic markers within domestic ELL populations, and how these markers subsequently impact automated or human writing evaluations. Because hedges and boosters anchor the two opposing poles of authorial certainty, this study intentionally isolates hedging, a decision also supported by practical computational constraints, to deliver a highly focused and technically efficient analysis of how language learners acquire nuanced pragmatic skills. Using the large-scale ELLIPSE corpus of argumentative essays written by U.S. students in grades 8–12 (Crossley et al., 2023), this study addresses these gaps through three primary research questions:

- **RQ1:** To what extent do demographic characteristics, including grade level, gender, race, and SES, predict variation in hedge use density among ELL student writers in the ELLIPSE corpus?
- **RQ2:** Are certain demographic groups more likely to use specific types of hedges (Epistemic Verbs, Epistemic Adverbs, Epistemic Adjectives, Epistemic Modal Verbs, and Approximation)?
- **RQ3:** Does hedge density predict overall writing quality and analytic subscale scores after accounting for essay-level linguistic features and student demographic background?

This study offers significant theoretical, empirical, and methodological contributions to the fields of learning analytics and educational linguistics. Methodologically, it advances the field by combining traditional corpus linguistics with an LLM-assisted hedge annotation pipeline and pilot human validation to efficiently identify nuanced, context-specific metadiscourse markers across a large dataset. Empirically, it provides a detailed investigation of how hedge density and hedge-type distributions vary across sociodemographic groups and relate to writing quality. More importantly, the study moves beyond large-scale annotation by examining whether these AI-derived discourse features function as interpretable indicators of rhetorical development rather than simple frequency-based measures.

3 Sample

3.1 Data Source and Variables

Data for this study come from the English Language Learner Insight, Proficiency and Skills Evaluation (ELLIPSE) corpus (Crossley et al., 2023), a publicly available dataset of 6,482 argumentative essays written by ELL students in grades 8 through 12 across the United States. The corpus includes rich demographic metadata (e.g., grade

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 416–422
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

level, gender, race, and SES), essay-level linguistic features and human-assigned holistic and analytic writing quality scores. Writing quality scores were assigned by trained raters on a scale of 1.0 to 5.0 across seven dimensions: Overall, Cohesion, Syntax, Vocabulary, Phraseology, Grammar, and Conventions. Additionally, essay-level linguistic variables include essay length (number of words), Type-Token Ratio (TTR, i.e., a bounded measure of linguistic diversity calculating the ratio of unique words to total words, which naturally decreases as text length increases) and Measure of Textual Lexical Diversity (MTLD, i.e., how diverse a writer's vocabulary is without being biased by the length of the text). The corpus spans grades 8 through 12, with the largest representation at grades 11 and 12. Full distributional details are in Figure 1.

4 Methodology

4.1 Metadiscourse Marker Detection for Hedge

To balance broad candidate coverage with contextual precision, hedge identification was conducted using a two-stage pipeline that combined dictionary-based candidate extraction with LLM validation across 6,482 essays. First, essays were segmented into sentences using NLTK tokenization, and candidate occurrences were flagged via a dictionary baseline drawn from a theoretically grounded taxonomy of five hedge categories: epistemic modal verbs, epistemic verbs, epistemic adverbs, epistemic adjectives, and approximations, including Hyland (2005), Aull and Lancaster (2014), Islam et al. (2020), Biber et al. (1999), Holmes (1988), and Hinkel (2004, 2005), with the full compiled list reported in Appendix A.

Second, because many candidate expressions are multifunctional, each occurrence was evaluated in its sentence context through prompt-based LLM validation. The prompt instructed the model to retain a candidate only when it reduced certainty, precision, or authorial commitment in the writer's own proposition. The model returned structured JSON outputs identifying whether the candidate functioned as a hedge and, if retained, its hedge category.

Prior to full-sample application, a pilot comparison with manual human coding was conducted to refine the prompt and examine whether LLM classifications aligned with the study's operational definition of hedging. This pilot served primarily as a formative validation step rather than a full reliability study; therefore, the resulting annotations should be interpreted as human-informed LLM classifications rather than fully established gold-standard labels. Future work will extend this validation using a larger independently coded sample and formal agreement metrics.

After applying this pipeline, 6,473 of 6,481 essays had at least one candidate, and of the 146,833 candidates, 58,811 were confirmed as actual hedges by the LLM (~40%), from which metrics such as total count and density per 100 words were extracted.

4.2 Statistical Analysis

Three sets of regression models were estimated to address the research questions. For RQ1, ordinary least squares (OLS) regression models were estimated with hedge density as separate outcome variables and demographic characteristics as predictors. Models were built using a stepwise entry approach, with text-level controls including word count, type-token ratio, and MTLD entered prior to demographic predictors, allowing the independent contribution of demographic variables to be assessed after accounting for essay-level linguistic features.

To address RQ2, we used multivariate analysis of variance (MANOVA) to examine whether demographic variables predicted students' hedge-type profiles across the five hedge categories. Category-specific hedge densities served as the multivariate outcome. Because MANOVA evaluates group differences across multiple dependent variables simultaneously, it was appropriate for testing whether demographic factors were associated with different distributions of hedge types rather than simply different total frequencies.

For RQ3, a series of nested OLS regression models were estimated predicting writing quality scores from hedge density respectively. Each marker was first entered alone as a predictor, then with essay-level text controls, and finally with the full set of demographic predictors added. This stepwise structure allows the persistence of the hedge effects across model specifications to be evaluated. Models were estimated separately for the overall holistic score and for each of the six analytic subscales: Cohesion, Syntax, Vocabulary, Phraseology, Grammar, and Conventions. All analyses were conducted in Python. Effect sizes are reported as unstandardized regression coefficients with standard errors, and statistical significance is evaluated at $\alpha = 0.05$.

5 Anticipated Results

Preliminary analyses suggest that overall hedge density is not strongly or consistently explained by students' broad demographic characteristics (Figure 2), where demographic factors (grade, gender, SES, race) account for only 0.9% of the total variance in hedge density. Descriptive comparisons showed relatively similar rates of hedging across gender, SES, grade level, and race groups (Figure 3), indicating that the total frequency of hedging may be shaped more by writing proficiency, lexical resources, prompt demands, genre familiarity, or individual rhetorical style than by demographic background alone. This pattern suggests that hedge density should be interpreted cautiously as an AI-derived writing indicator: higher or lower hedge frequency does not necessarily correspond to stronger or weaker writing, nor should it be treated as a simple demographic marker.

However, the distribution of hedge types appears to provide a more sensitive measure of students' rhetorical development. Students relied most heavily on epistemic modal verbs, such as *may*, *might*,

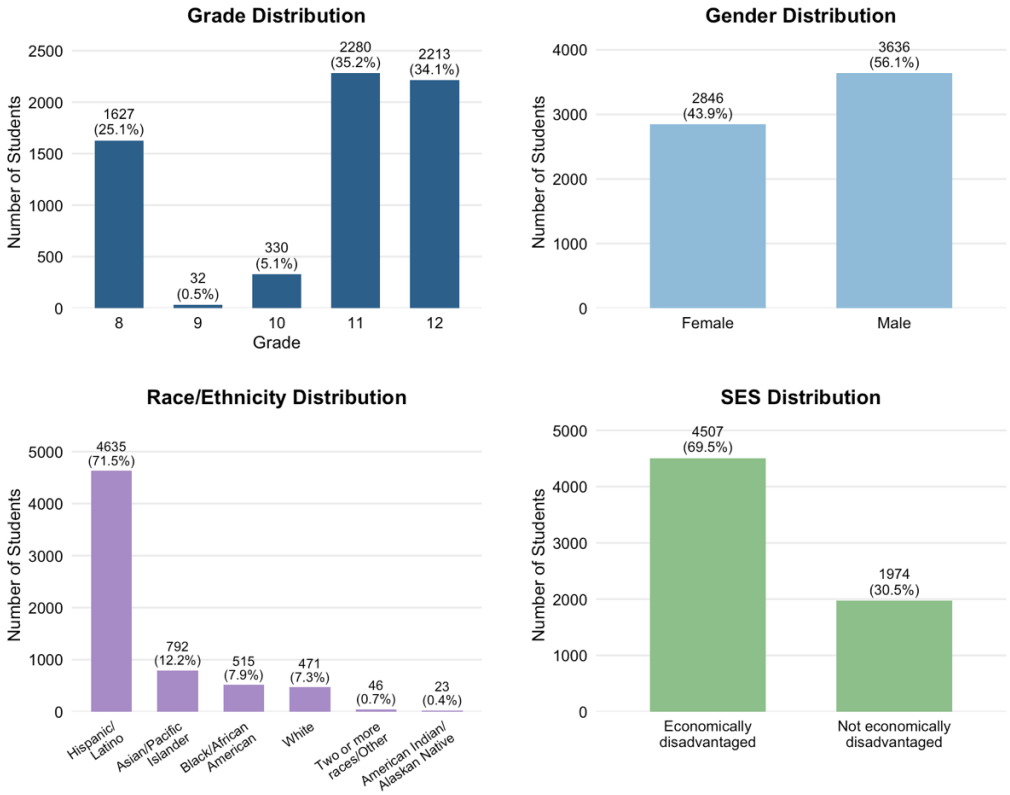


Figure 1: ELLIPSE corpus sample characteristics (N = 6,482).

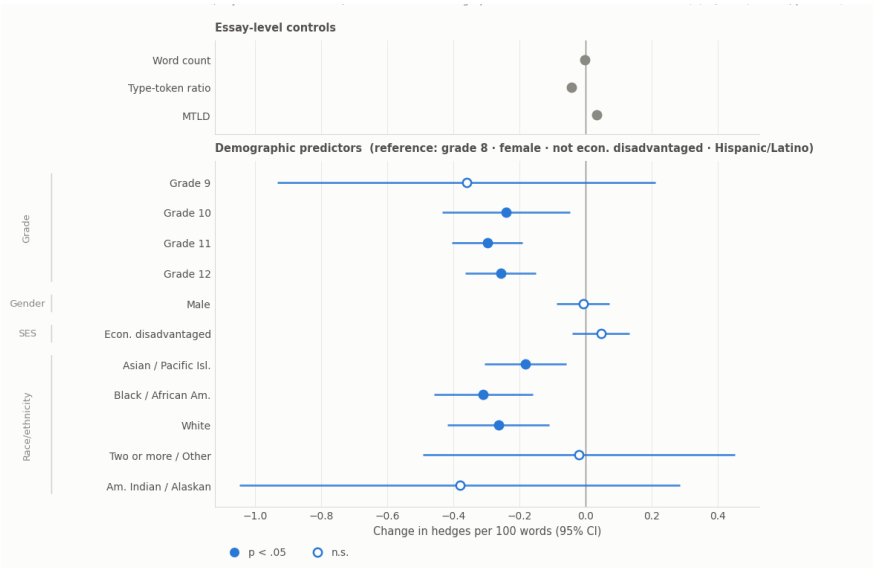


Figure 2: OLS regression coefficients for demographic predictors of overall hedge density (adjusted for text-level controls).

can, *could*, and *would*, and on approximation markers, such as *some*, *many*, *most*, *about*, and *often*, which accounts for 83% of all hedge tokens in total (Figure 4). In contrast, more lexically specific stance markers, including epistemic verbs, epistemic adverbs, and epistemic adjectives, appeared less frequently. This pattern suggests that many ELL writers may have access to basic hedging strategies but may use a relatively narrow repertoire of stance-marking resources.

The MANOVA results further suggest that grade level may be associated with students' hedge-type profiles even when overall hedge density remains relatively stable. Figure 5 shows that grade level has the largest multivariate effect size (partial $\eta^2 = 0.0087$, $F = 11.30$, $p < 0.001$) compared to all other demographics, though all effects are significant. In other words, developmental differences may appear less as a simple increase or decrease in the amount of hedging and more as a shift in the kinds of hedges students use. Follow-up analyses in Figure 6 point especially to differences in

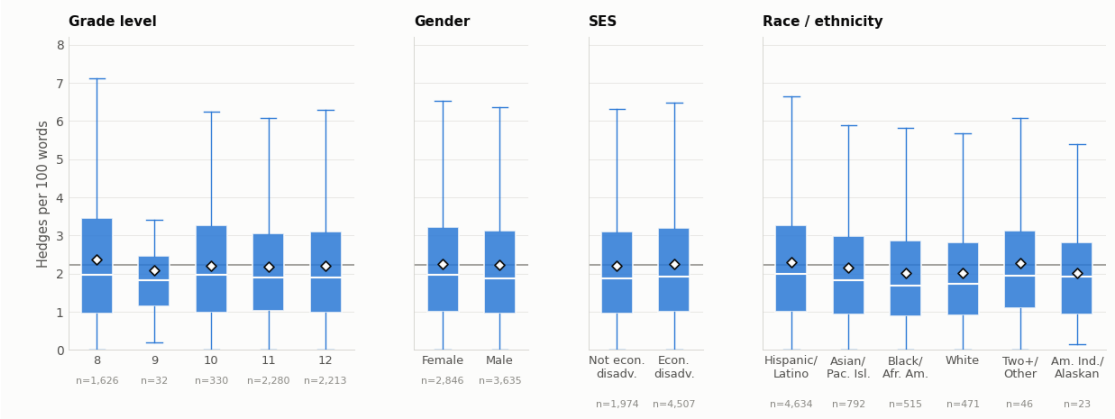


Figure 3: Distribution of total hedge density by group.

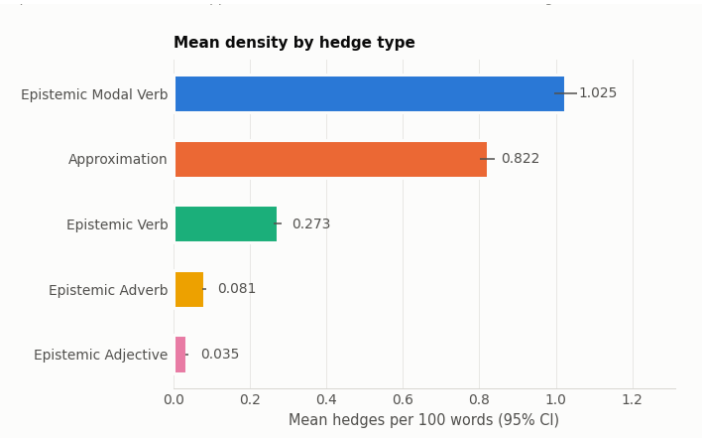


Figure 4: Mean hedges per 100 words by hedge type.

modal verbs and approximation markers, suggesting that students across grade levels may vary in how they balance familiar, high-frequency forms of uncertainty marking with more precise or lexically nuanced forms of epistemic stance.

Findings for gender, SES, and race were substantially weaker and should be interpreted cautiously. While some model-based estimates suggest possible group differences in hedge density or category use, these effects are not consistent across all descriptive and inferential tests (Figures 7 and 8). Rather than treating demographic differences as evidence of deficit, the study interprets them as potential indicators of varied rhetorical socialization, language-learning histories, and access to academic writing resources.

For writing quality, we anticipate that hedge density alone will have limited explanatory power after controlling for essay length, lexical diversity, and demographic background. If hedge measures remain associated with holistic or analytic writing scores, this would provide evidence that AI-derived metadiscourse indicators capture a meaningful aspect of argumentative writing. However, the expected pattern is not that "more hedging is better." Instead, writing quality is likely to depend on whether students use hedges flexibly and appropriately to qualify claims, align certainty with evidence, and construct a credible academic stance.

Overall, the anticipated results suggest that hedge-type distribution and hedge variety may be more instructionally and psychometrically useful than raw hedge frequency. For AI-integrated assessment, this means that automated metadiscourse indicators should not be interpreted as simple counts. Instead, they should be evaluated as evidence of students' developing rhetorical range, pragmatic control, and ability to manage certainty in academic argumentation.

6 Anticipated Conclusions & Contributions

If supported by the full analyses, the findings suggest that epistemic stance should receive more explicit attention in ELL writing instruction. Students may need support not only with vocabulary and grammar, but also with how to express uncertainty, qualification, and evidence-based confidence in academic arguments. Instruction should avoid simply encouraging students to use more or fewer hedges. Instead, teachers can help students select stance markers according to evidence strength and rhetorical purpose. For example, students can compare expressions such as "This proves that," "This suggests that," and "This may indicate that" to examine different degrees of certainty.

For assessment and learning analytics, the study suggests that hedge density alone should not be treated as a direct indicator of writing quality. If category-specific patterns or hedge variety remain

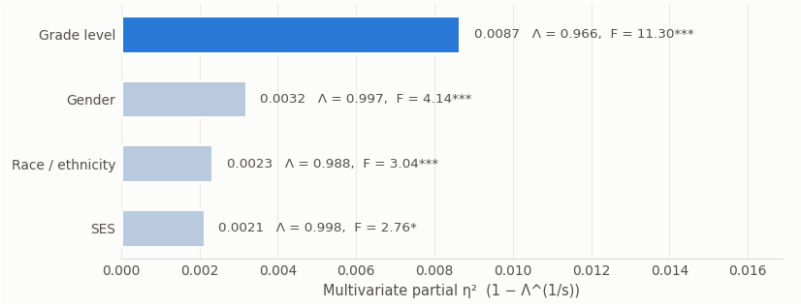


Figure 5: MANOVA on the hedge type. Note that *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

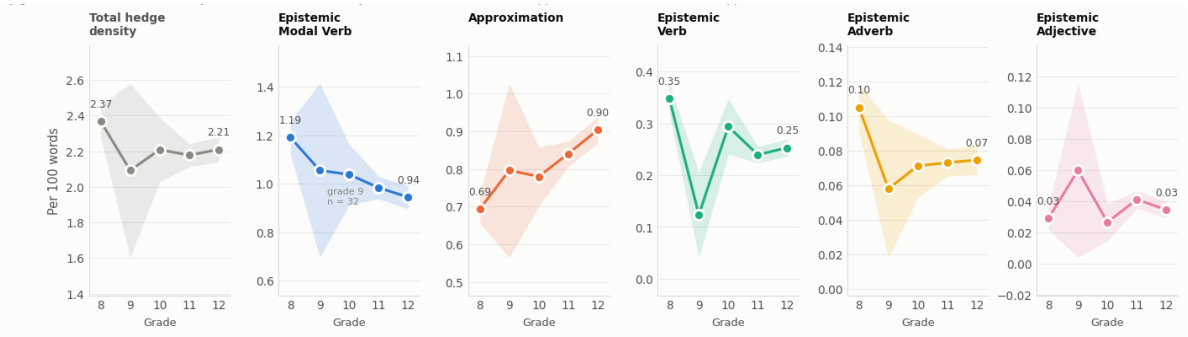


Figure 6: Mean density of total hedges and specific hedge categories across Grades 8–12.

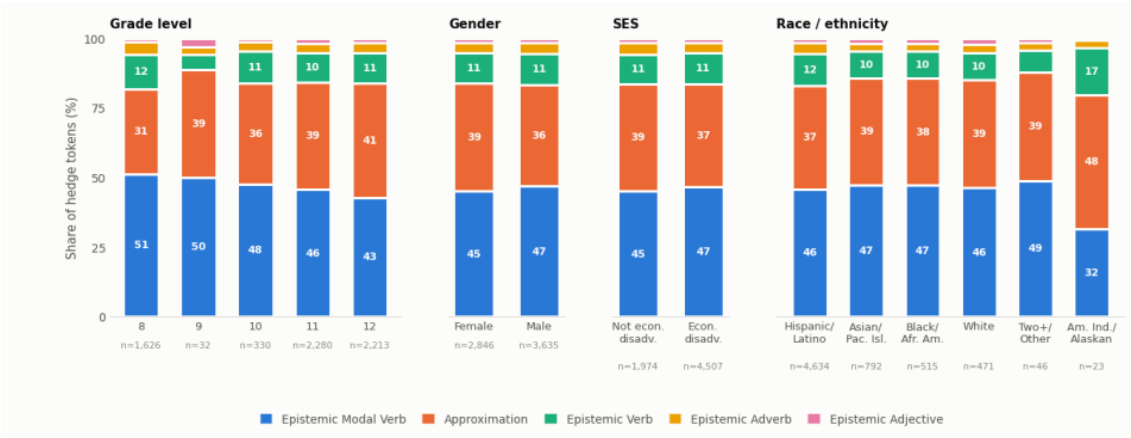


Figure 7: Proportional distribution of hedge categories across demographic subgroups.



Figure 8: Univariate Type II ANOVA F-statistics for demographic predictors across hedge categories. Darker shading indicates larger test statistics; cells marked n.s. denote non-significant effects ($p \geq 0.05$).

associated with writing performance after linguistic controls are introduced, these features may provide more interpretable evidence of students' rhetorical repertoire than raw frequency alone. Learning

analytics tools could therefore support formative feedback by identifying overreliance on a limited set of stance markers and encouraging a broader range of context-appropriate expressions.

The demographic analyses should also be interpreted cautiously. Differences across grade, gender, race, or SES do not in themselves demonstrate assessment bias or fairness. Rather, they provide fairness-relevant evidence by showing whether AI-derived discourse indicators vary systematically across student groups and by identifying patterns that may call for further investigation in future validation studies. In this sense, the present study contributes to fairness research by informing how such indicators should be interpreted, rather than by establishing that the assessment system itself is fair or biased.

Limitations

A key limitation of this study is its exclusive focus on hedges rather than examining both hedges and boosters. Because these features represent different forms of stance calibration, future work should examine a broader range of metadiscourse markers to provide a more complete account of how ELL students construct academic stances.

References

A. M. Ahmed and X. Zhang. 2023. Students' voice in L2 English writing: A systematic review of literature. *Ampersand*, 10:100114. doi:10.1016/j.amper.2023.100114.

S. Alanazi, A. H. Al-Hoorie, and A. Alghamdi. 2024. The use of interactional metadiscourse markers by Saudi EFL male and female college students: The case of a gender-sensitive topic. *Humanities and Social Sciences Communications*, 11(1):988. doi:10.1057/s41599-024-03506-3.

L. L. Aull and Z. Lancaster. 2014. Linguistic markers of stance in early and advanced academic writing: A corpus-based comparison. *Written Communication*, 31(2):151–183.

C. Bazerman. 2009. Genre and cognitive development: Beyond writing to learn. *Pratiques. Linguistique, littérature, didactique*, 143–144:127–138.

D. Biber. 2004. Historical patterns for the grammatical marking of stance: A cross-register comparison. *Journal of Historical Pragmatics*, 5(1):107–136.

D. Biber, S. Johansson, G. Leech, S. Conrad, and E. Finegan. 1999. *Longman Grammar of Spoken and Written English*. Longman.

C. Castillo and N. Ávila Reyes. 2026. Students' sociodemographic characteristics and writing performance: A systematic literature review. *Reading and Writing*, 39(6):1085–1121. doi:10.1007/s11145-025-10658-4.

S. Crossley, Y. Tian, P. Baffour, A. Franklin, Y. Kim, W. Morris, and U. Boser. 2023. The English language learner insight, proficiency and skills evaluation (ELLIPSE) corpus. *International Journal of Learner Corpus Research*, 9(2):248–269.

D. A. S. El-Dakhs. 2020. Why are hedges used differently in research articles and opinion columns? A contrastive genre analysis. *Journal of Pragmatics*, 156:63–75. doi:10.1016/j.pragma.2019.09.009.

M. Fuoli, W. Huang, J. Littlemore, S. Turner, and E. Wilding. 2025. Metaphor identification using Large Language Models. *arXiv preprint arXiv:2509.24866*.

E. Hinkel. 2004. *Teaching Academic ESL Writing: Practical Techniques in Vocabulary and Grammar*. Lawrence Erlbaum Associates.

E. Hinkel. 2005. Hedging, inflating, and persuading in L2 academic writing. *Applied Language Learning*, 15(1/2):29.

J. Holmes. 1988. Doubt and certainty in ESL textbooks. *Applied Linguistics*, 9(1):20–44.

K. Hyland. 1998. *Hedging in Scientific Research Articles*. John Benjamins.

K. Hyland. 2005. *Metadiscourse: Exploring Interaction in Writing*. Continuum.

J. Islam, L. Xiao, and R. E. Mercer. 2020. A lexicon-based approach for detecting hedges in informal text. In *Proceedings of the Twelfth Language Resources and Evaluation Conference (LREC '20)*, pages 3109–3113. European Language Resources Association.

J. J. Lee and L. Deakin. 2016. Interactions in L1 and L2 undergraduate student writing: Interactional metadiscourse in successful and less-successful argumentative essays. *Journal of Second Language Writing*, 33:21–34. doi:10.1016/j.jslw.2016.06.004.

G. Lorenz. 1999. *Adjective Intensification—Learners versus Native Speakers: A Corpus Study of Argumentative Writing*. Rodopi.

E. Ochs and B. B. Schieffelin. 2011. The theory of language socialization. In *The Handbook of Language Socialization*, pages 1–21.

S. Park and S. Y. Oh. 2018. Korean EFL learners' metadiscourse use as an index of L2 writing proficiency.
SNU Journal of Education Research, 27(2):65–89.

W. S. Pearson and E. Abdollahzadeh. 2023. Metadiscourse in academic writing: A systematic review. *Lingua*, 293:103561.

P. Uccelli, C. L. Dobbs, and J. Scott. 2013. Mastering academic language: Organization and stance in the persuasive writing of high school students. *Written Communication*, 30(1):36–62.

M. T. Yeganeh and S. M. Ghoreyshi. 2015. Exploring gender differences in the use of discourse markers in Iranian academic research articles. *Procedia - Social and Behavioral Sciences*, 192:684–689.

Appendix A. Hedge Taxonomy

Hedge Category	Definition	Examples	Research Evidence
Epistemic Verb	A lexical verb that presents a proposition as inferred, believed, assumed, suggested, or only partially certain, rather than fully asserted as fact. Do not count a candidate word as a hedge if it only refers to the content of thinking, believing, or opinion-holding, rather than reducing the writer's own commitment in context. Do not count *think* as a hedge when it only refers to the content of someone's thoughts or beliefs, rather than softening the writer's own assertion.	suggest, believe, appear, indicate, assume, seem, consider, doubt, estimate, expect, feel, guess, imagine, speculate, suppose, think, understand, imply, presume, suspect, postulate, reckon, infer, hope	Biber treats verbs that control stance/complement constructions as a major way English expresses stance, and Varttala-style hedge taxonomies treat these as a core hedge class.
Epistemic Adverb	An adverb that modifies a clause or proposition to signal likelihood, tentativeness, limited precision, or non-absolute commitment.	rather, slightly, barely, strictly, presumably, fairly, theoretically, basically, relatively, possibly, preferably, slenderly, scantily, decidedly, arguably, seemingly, occasionally, partially, partly, practically, roughly, virtually, allegedly	Adverbials are a major grammatical resource for expressing epistemic meaning; Biber distinguishes stance adverbials as a major type of stance marking.
Epistemic Adjective	An adjective that encodes possibility, probability, uncertainty, or appearance.	presumable, possible, probably, likely, apparent, probable, improbable, unlikely, rarely, improbably, unclearly, unsure, sure, chance, unclear	Adjectives like *possible* and *likely* as expressing modal-like meanings, and hedge classifications regularly include epistemic adjectives as a core category.
Epistemic Modal Verb (Modal Auxiliaries)	A modal auxiliary used to indicate possibility, weak prediction, or non-categorical commitment rather than certainty. Count modal verbs as hedges only when they express uncertainty, possibility, or limited commitment, not when they express obligation, advice, or normative judgment.	may, might, maybe, shall, should, can, could, would, ought	Modal auxiliaries are one of the most established hedge types, namely can, could, may, might, must, should, will, would.
Approximation	A word or phrase that makes a claim less exact, less categorical, or more general by signaling approximate quantity, frequency, degree, or scope.	usually, approximately, generally, frequently, normally, some, many, most	Hyland's line of work treats accuracy-related hedging as reducing precision, Salager-Meyer's classic taxonomy includes approximators, and Hinkel includes frequency and quantity softeners such as normally, relatively, often, frequently, and usually among hedge-like resources in academic writing.

Beyond the Grid: Auditing Code-Editing Search for Automated Essay Scoring

Julius Frost
Cambium Assessment
julius.frost@cambiumassessment.com

Abstract

Code-editing language-model agents can change training choices beyond the hyperparameter grids tested in automated essay scoring. Comparing these approaches requires distinguishing gains from a broader search space from evidence of a better search procedure. We compare code-editing agents with grid-restricted search, including random search, in two studies on ASAP-AES with nominally matched 12-hour search budgets. Code-editing produced the configuration with the highest test quadratic-weighted kappa (QWK) point estimate in each primary comparison. However, random search over a grid built afterwards around the Study 2 agent's backbone, input length and head rule recovered most of its gain over BERT. Requiring a minimum validation-score improvement to accept a trial also left the accepted configuration's validation score below the highest recorded valid validation score in every run that accepted a trial under this rule. The primary comparisons use one search run per condition, and some test folds reused essays involved in configuration selection, so these results compare selected configurations without establishing search-procedure superiority. These findings motivate reporting the available search choices and both accepted and best-scoring trials, and evaluating search procedures through repeated runs on data independent of configuration selection.

1 Introduction

Comparisons of code-editing search in automated essay scoring depend on the harness: its editable search space, acceptance rule, and evaluation procedure determine what an agent can propose and what the comparison can establish. We compare the test QWK of configurations selected by code-editing and grid-restricted search under nominally matched 12-hour budgets, and examine what acceptance and evaluation records allow us to conclude. Our automated essay scoring (AES) base grids fix the backbone and loss; hyperparameter optimization (HPO) can also search these choices. Large language models (LLMs) can edit training code. Karpathy's `autoresearch` repository[1] set out the pattern: a single editable training file, a protected evaluator with one metric, and an accept-or-reject decision per trial. Ferreira et al. (2026) compared fixed-space HPO and LLM search, including source editing, on language-model pretraining.

Essay scoring's ordinal objective gives code-editing search choices beyond numeric hyperparameters: the loss, head, sampler, and score discretization can all change how the model learns or predicts scores. We compare search conditions on ASAP-AES from two starting configurations: one untuned to this dataset, the other adapted from an AES literature baseline. Screening selects configurations using validation scores on a prompt subset; we then evaluate them across prompts and folds.

Contributions. First, the code-editing configuration has the highest test-QWK point estimate in each primary comparison (Table 1). Random search over Study 2's grid, built afterwards around the code-editing configuration's backbone, input length and head rule, recovers most of that configuration's gain over BERT. Second, requiring a minimum improvement margin left accepted trials below the highest valid screening QWK in every run that accepted a trial under this rule; we report both recorded scores and trial identities (§4.2). Third, the primary comparisons use one search run per condition; on screened prompts, other folds' test splits reuse selection essays. The comparisons therefore concern selected configu-

[1] https://github.com/karpathy/autoresearch

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 423–440
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

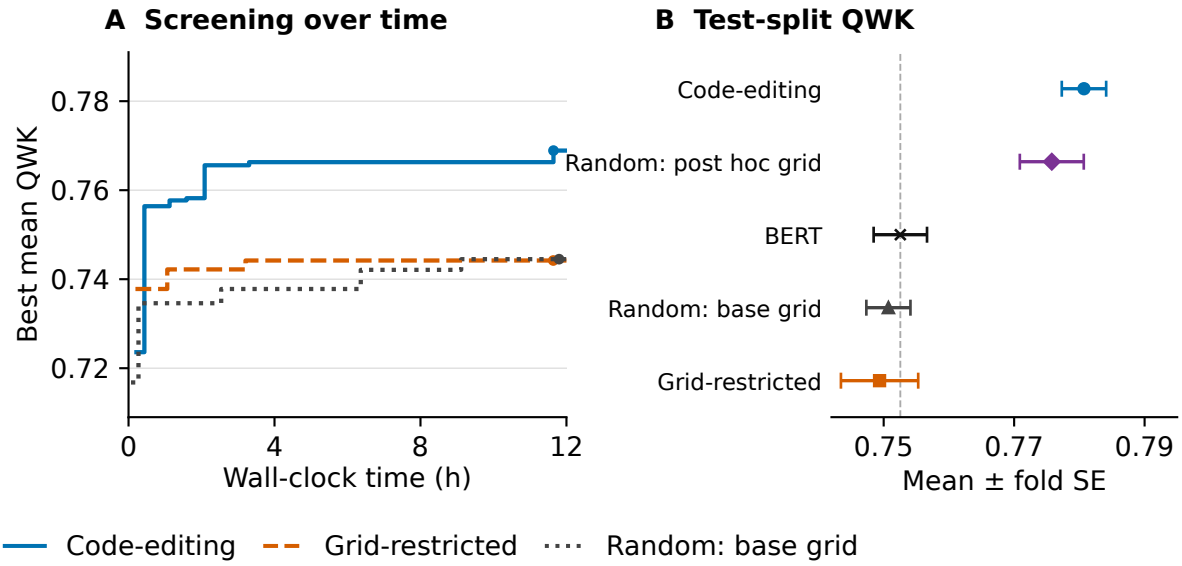


Figure 1: Code-editing has the highest test-QWK point estimate in Study 2's primary comparison; its fold-SE interval overlaps that of post hoc random search (Table 1). A: single-seed best screening QWK with unequal starting scores; condition-specific time origins preclude speed comparisons (§4.3.2). B: corrected test QWK, mean ± fold standard error (SE) over 40 prompt-fold cells; dashed line, BERT. The code-editing result evaluates its accepted trial. The post hoc grid was built afterwards around that trial's backbone, input length and head rule. Each search ran once; screened-prompt test folds reuse selection essays.

rations (§3.1.1; §3.4), not search-procedure performance. These findings motivate reporting permitted search choices and both accepted and best-scoring trials; they do not establish which harness design improves search quality.

2 Related Work

Agents that do machine-learning work. MLAgentBench (Huang et al., 2024) evaluates LLM agents on end-to-end ML experimentation tasks and MLE-bench (Chan et al., 2025) scores them against Kaggle leaderboards. In AIDE (Jiang et al., 2025) the agent is a tree search over solution scripts; Toledo et al. (2025) model such agents as search policies with explicit operator sets, comparing greedy, Monte Carlo tree search and evolutionary policies on MLE-bench lite. These works evaluate agents through task outcomes or search-policy comparisons. We compare code-editing search with grid-restricted HPO under nominally matched time budgets, examining how edit permissions and acceptance rules affect the comparison.

LLMs as optimizers. OPRO (Yang et al., 2024) uses an LLM to propose candidates for black-box optimization; Liu et al. (2025) apply the idea to hyperparameters, and OPT-BENCH (Li et al., 2026) benchmarks agents on large search spaces. Code modification lets GrowthHacker (Wu et al., 2026) tune off-policy evaluators; van Stein et al. (2025) separate LLM-generated heuristic code from in-loop HPO. Closest to us, Ferreira et al. (2026) compared LLM agents with covariance matrix adaptation evolution strategy (CMA-ES) and the tree-structured Parzen estimator (TPE) inside Karpathy's loop on a pretraining loss. The HPO methods performed better on a fixed hyperparameter space. Source editing narrowed the gap without closing it; an LLM–CMA-ES hybrid sharing the optimizer's search state had the best mean. Their comparison clock excludes LLM inference, unlike ours. We examine ordinal scoring and use a post hoc grid to assess how much of the selected configuration's gain random search recovers. We also distinguish accepted trials from those with the highest screening scores.

Li (2026) reports no benefit from simulated annealing over greedy hill climbing in two tested tasks with LLM proposals. Our acceptance-rule comparison (§3) reports individual search-run outcomes, not an estimate of policy equivalence (§5).

Baselines and task. Random search (Bergstra and Bengio, 2012), the strongest tested HPO comparator in Study 1, supplies Study 2's fixed-grid comparator and post hoc diagnostic. The selected

ordinal losses belong to established surrogate-loss families (Hou et al., 2017; de la Torre et al., 2018). Our AES model and loss choices also overlap prior regression-and-ranking work (Yang et al., 2020; Xie et al., 2022); those ingredients could inform a prospective categorical baseline. Our AES baselines follow Ormerod and Kwako (2024) and the BERT configuration of Rodriguez et al. (2019); splits follow Taghipour and Ng (2016). Ormerod and Kehat (2025) study long-context AES, including ModernBERT, on ASAP 2.0 rather than our eight-prompt corpus.

3 Methods

3.1 Shared protocol

A *harness* is the execution, scoring, and budget machinery surrounding the agent. The *autoresearch loop* is the agent's propose, edit, evaluate, and accept-or-reject cycle. A *condition* specifies the compared search method, edit permissions, acceptance rule, and budget; a *search run* is one budgeted execution of that condition. A *trial* is a candidate's training/evaluation attempt, while a *configuration* fixes the training choices. Trial counts need not equal the number of distinct configurations (Appendix A). *Screening* selects configurations using validation data on a prompt subset. *Post-search retraining and test-split evaluation* evaluates a selected configuration on rotating folds; it does not supply an outer holdout for search (§3.1.1).

Our harness retains Karpathy's workflow (Appendix H). Each original search run has a 12-hour wall-clock budget, including thinking and trials. Study 1 fixes epochs, input length and data fraction except for halving (§3.2.1); Study 2 permits configuration changes. The agent selects on validation; post-search runs score selected configurations on test splits. In Study 2 the launcher also removes read permission on test files. This within-run access rule does not establish independence between selection and test data across folds.

We replace Karpathy's any-improvement rule with a requirement that a trial's screening QWK exceed the incumbent score (the accepted configuration's screening QWK) plus a minimum improvement margin estimated from baseline training-seed reruns (Appendix G). This sum is the acceptance threshold; equality fails. The margin is not a significance threshold. An accepted trial becomes the starting configuration for the next proposal. It need not have the highest valid screening QWK: better-scoring trials can fall short of the acceptance threshold.

Trials are recorded in an experiment tracker. The agent instructions also require a git commit stating each proposed change's rationale. A trial counts only if it trained, evaluated the required split at the required fidelity, and covered all four screening prompts (§3.1.1). The agent instructions (Appendix H) name the metric, files, minimum improvement margin and exit-code contract, and state that no additional dataset-specific knowledge is supplied. They retain Karpathy's preference for simpler code and improvements obtained by deleting code. The agent otherwise receives a single launch line.

3.1.1 Testbed and metric

ASAP-AES has eight essay prompts with reported dataset score ranges from 0–3 (prompt 3) to 10–60 (prompt 8) (Ormerod and Kwako, 2024). We use supplied local split files following the five-fold 60/20/20 protocol of Taghipour and Ng (2016), scoring each prompt by QWK on the real score scale and averaging the eight prompt QWKs with equal weight (macro-averaging). In both studies the screening objective is the mean validation QWK across four prompts {2, 3, 5, 8} spanning the reported score-range extremes. Post-search evaluation retrains selected configurations on each fold's training split, selects an epoch on validation, and scores its test split. The configuration, not trained weights, is reused across eight prompts $\times$ five folds (40 prompt-fold cells per condition). Prompts {1, 4, 6, 7} were not screened. On screened prompts, rotating test folds include essays from the selection population, so the aggregate is not a fully independent test of search.

3.2 Editable search space

3.2.1 Study 1: configuration not tuned to ASAP-AES

Study 1 starts from an Ettin/LoRA configuration that had not been tuned to ASAP-AES before the benchmark, unlike Study 2's dataset-informed literature baseline. All five conditions train an Ettin-encoder-400m model (Weller et al., 2026) with low-rank adaptation (LoRA; Hu et al., 2022) through the shared harness with the same screening objective. Screening runs at reduced fidelity (4 epochs, max length 640, half the training data, one fold, one seed), under the shared search budget.

Grid search, random search, TPE (Bergstra et al., 2011), and halving search a fixed 864-point grid over seven LoRA and optimizer axes (Appendix F) with the starting cross-entropy `train.py`. The code-editing condition (Claude Opus 4.8, unattended, logged as `autoresearch`) may set the same grid values and edit five designated functions: loss, head, sampler, training-data transform, and score discretization. The minimum improvement margin is 0.02; its empirical basis is recorded in Appendix G. Each condition's best screening configuration is retrained at full fidelity (10 epochs, max length 1536, all training data) and evaluated on test splits for all 40 cells. The condition logged as asynchronous successive halving (ASHA; Li et al., 2020) is labeled halving: the supplied scheduler uses synchronous successive halving (Jamieson and Talwalkar, 2016). Counts and the unauthenticated historical scheduler are described in Appendix A.

3.2.2 Study 2: AES literature baseline

The starting configuration adapts Rodriguez et al.'s BERT classifier (Rodriguez et al., 2019; Devlin et al., 2019). Implementation settings are in Appendix K.

Two conditions run Claude Sonnet 5 (Claude Code 2.1.261), each from the same baseline on its own L40S under the shared budget. They differ in edit permissions. The **code-editing** condition (GPU A) may edit `train.py` outside its protected `main()` entry point; the **grid-restricted** condition (GPU B) may only pass a configuration drawn from a 192-point grid over six optimizer axes (Appendix F), with `train.py` hash-pinned (checksum-verified each trial) and the driver refusing any other change. Each condition uses its own GPU's baseline screening score. The minimum improvement margin is 0.036 (Appendix G). After search, accepted, highest-screening, and selected diagnostic configurations are retrained for evaluation on all 40 cells. Search-run counts and screening scores retain their historical values; the later audit does not revise them.

Two no-agent random-search conditions use the same budget: the base numeric grid and a post hoc grid built after observing the accepted configuration's backbone, input-length, and head-rule choices. Prompt-dependent head selection uses classification for at most ten training-observed score levels and sigmoid regression otherwise. The post hoc grid restricts numeric choices; no prospective categorical comparator was tested (Appendix F; Table 1).

3.3 Acceptance rule

The acceptance-rule ablation replaces the minimum improvement margin with any-improvement acceptance; its agent instructions otherwise match the Sonnet code-editing condition. Any-improvement acceptance requires screening QWK strictly above the incumbent score (§3). We distinguish reclassifying a recorded trajectory under a different rule from executing a search run with that rule (Table 2).

3.4 Evaluation outcomes

Selection in both studies is on QWK only. Yannakoudakis and Cummins (2015) caution against kappa because of its sensitivity to score distributions and scoring scales. We report six secondary outcomes for Study 2 (Table 9) and all except standardized mean difference for Study 1 (Table 8). The agreement and distribution conventions follow the operational-scoring framework of Williamson et al. (2012). Exact and adjacent (± 1) agreement measure the share of essays scored exactly right or within one point (0–1). Mean absolute error (MAE) is in raw score points. Standardized mean difference (SMD) is machine minus human mean divided by the square root of their mean variance (0 at parity). The predicted-to-human score standard deviation (SD) ratio is 1 at parity. Exact agreement on extreme-score essays uses each prompt/fold's observed minimum and maximum human scores, weighted by essay count across those endpoints (0–1).

Fold standard errors describe variation conditional on a configuration; fold dependence complicates variance estimation (Bengio and Grandvalet, 2004). Shared data and overlapping training folds preclude treating folds as independent search replicates. We report point differences, not significance or equivalence tests. Study 1's QWK SE propagates per-prompt population-SD estimates; Study 2's uses sample SD of five fold means divided by $\sqrt{5}$ (Appendix G).

3.4.1 Evaluation audit

For Study 2 we compute QWK from every prediction row on the contiguous integer score scale; the other six outcomes also include all rows. The audit compares these scores with the historical logs. `study2_metrics.py` reproduces

Study	Condition / evaluated configuration	Test QWK	SE
1: untuned Ettin/LoRA	code-editing (LLM), highest screening	0.771	0.0046
	random	0.752	0.0041
	halving	0.739	0.0073
	grid	0.715	0.0073
	TPE	0.654	0.0301
2: adapted BERT	BERT-base comparator	0.753	0.0041
	code-editing: str008 (accepted)	0.781	0.0034
	grid-restricted: num020 (highest-screening)	0.749	0.0059
	random, base grid	0.751	0.0034
	random, post hoc grid	0.776	0.0049
Derived contrast		QWK difference / share	SE
Study 1: code-editing minus random test QWK		+0.019	≈ 0.006
Study 2: post hoc random gain over BERT		+0.023	0.0059
Study 2: str008 residual over post hoc random		+0.005	0.0042
Study 2: post hoc share of str008 gain		82%	–

Table 1: Search space comparison (40 prompt-fold cells). SE: logged propagation (Study 1), sample fold SE (Study 2). Contrasts: quadrature ignoring covariance (Study 1), paired folds (Study 2; Appendix G). Share: unrounded BERT-relative gain ratio; dash: not applicable.

corrected means and contrasts without overwriting those logs (Appendix G; audit findings in §4.3.2). An outer holdout or nested search would be needed to assess the search procedure independently (Cawley and Talbot, 2010).

4 Results

4.1 Editable search space

4.1.1 Study 1: configuration not tuned to ASAP-AES

The code-editing condition had the highest screening and full-fidelity test QWK; random search was the strongest tested HPO comparator (Table 1; Appendix A). TPE's screening rank reversed at full fidelity, with prompt-level declines (Table 13). The grid condition covered only part of its specified space; each condition is one search run.

The code-editing advantage over random search concentrates on prompts 7 and 8; the other six prompts' mean contrast is near zero (Table 13). Study 1 predictions are unavailable to check the extreme-score omission found on these prompts in Study 2 (§4.3.2). The recorded configuration replaced cross-entropy with squared Earth Mover's Distance loss and added a square-root inverse-frequency sampler (Table 15). The accepted trials changed the loss or sampler (Appendix I). The evaluated configuration, trial 005, doubled the last accepted configuration's learning rate. It had the highest valid screening QWK but failed the acceptance threshold. The trial log associates the gains with these changes but does not isolate the loss or sampler's contribution.

4.1.2 Study 2: AES literature baseline

The code-editing condition accepted a configuration beyond the base grid, while the grid-restricted condition accepted none (Tables 2 and 4). The accepted configuration combines ModernBERT-base (Warner et al., 2025), longer input windows, and prompt-dependent head selection (Table 16). Relative to BERT, the accepted configuration's largest prompt-level test-QWK gain is on prompt 8; test QWK also rises on prompts 1 and 7, which were never screened (Table 14). The rejected diagnostic str001 also exceeds the BERT baseline (Appendix A).

Base-grid random search and the grid-restricted LLM both score below BERT; close point estimates do not establish equivalence. Post hoc random search recovers most of str008's BERT-relative gain (Table 1). The comparison is retrospective, with different evaluated epoch counts and an unverified executed head threshold (Appendix K). The residual test-QWK difference is smaller than str008's range across nominal-seed retrainings; these retrainings do not replicate search (Appendix B). Evaluation did not match configurations to GPU; the baseline offset remains uncontrolled (§4.3.2). Related post hoc diagnostics tune intermediate programs' constants (Pelleriti et al., 2026); neither that comparison nor ours establishes that iterative discovery was necessary.

The additional code-editing search run accepted no trial, although its best trial exceeded BERT in evaluation (Appendix B). On the common ModernBERT backbone, component con-

Acceptance record	Minimum improvement margin	Any improvement
Trials / accepted	45 / 1	51 / 5
Baseline screening QWK	0.724	0.741
Last accepted screening QWK	0.766	0.772
Best screening QWK	0.769	0.772
Last accepted configuration	str008	k042
Last accepted test QWK (corrected)	0.781	0.776
Sample fold SE	0.0034	0.0028
Test QWK gain over BERT	+0.028	+0.023

Table 2: Acceptance-rule comparison. QWK: quadratic-weighted kappa. Screening scores are historical records. Corrected test outcomes use 40 prompt-fold cells per last accepted configuration; SE describes fold variation conditional on that configuration (Appendix G).

trasts favored head changes over input length (Appendix D). These are conditional contrasts; they do not isolate the BERT-relative gain.

4.2 Acceptance rule

In Study 2's Sonnet code-editing run, the margin rule accepted an early trial but rejected later improvements (Appendix G). Applying any-improvement acceptance to that recorded trajectory would accept str012 and str044; it does not reconstruct proposals under a different acceptance history.

The separately executed acceptance-rule ablation accepted more trials than the original Sonnet margin-rule run, but none of its recorded scores clears its initial baseline plus the original margin (Table 2). Only its last accepted configuration was evaluated. Its recovered source combines ModernBERT-base, longer input windows, and scalar regression with pairwise ranking loss (Table 16). Ranking has precedent in text assessment (Yannakoudakis et al., 2011); the exact historical execution remains unverified. Its final trials again adjust numeric constants. The ablation holds the clock and remaining agent instructions fixed but does not isolate the margin's effect on later proposals or search quality. The other margin-rule runs that accepted trials are reported in Appendix I (Study 1) and Appendix C (Study 2 Opus).

4.3 Evaluation outcomes

4.3.1 Secondary scoring outcomes

The Study 1 LLM configuration has higher adjacent agreement and lower mean absolute error than random, halving, grid, and TPE, but lower exact agreement than grid and lower extreme-score agreement than random (Table 8). These comparisons do not isolate the cause of the tradeoff.

Relative to the BERT baseline in Study 2, the accepted code-editing configuration has higher adjacent and extreme-score agreement and lower mean absolute error (Table 9). Its SD ratio moves toward parity while SMD moves away (§3.4); exact agreement declines. The acceptance-rule ablation's k042 also loses exact agreement relative to BERT despite improving QWK (Tables 9 and 2). Appendix E retains the reported exact-agreement contrasts and uncertainty; these estimates exclude search-level variation and do not establish equivalence.

4.3.2 Evaluation audit

Elapsed times in Figure 1A include search overhead and are measured from the recorded agent launch for the Sonnet conditions and from the cumulative sweep timer for random search. The original agent budget-clock origins are unavailable, so the trajectories do not support precise equal-budget timing comparisons.

The historical evaluators omit one extreme-score essay in each affected prompt/fold cell by restricting labels to the training-observed range (Table 11). The recovered so042 source uses test-fold essays as unlabeled pretraining text on three folds; its transductive diagnostic is reported separately and excluded from cross-condition rankings (Appendix K). Corrected Study 2 scores include every prediction row, but validation predictions are unavailable: the correction has not been shown to preserve epoch selection, trial ranking, or acceptance decisions. Study 1 predictions are unavailable for this check, so its QWK remains log-derived.

The logged and corrected BERT means differ from the published comparison (Tables 11 and 1); the records do not establish identical implementation. The between-GPU screening off-

set is within the baseline seed spread but its cause is not isolated. Against the higher grid-restricted baseline, the accepted code-editing trial would miss the margin (Table 2; Appendix G). Corrected QWK matches an independent all-essay calculation on the contiguous integer score scale for every manifest-listed cell (Table 11). Historical-versus-corrected discrepancies exceed four-decimal rounding.

Historical split assignments, information access and executed-code identities remain unauthenticated (Appendix K). The recovered starting script creates the model before seeding; nominal seed pairing therefore does not establish identical initialization.

5 Discussion

Search-space specification. The selected loss/sampler and backbone/input/head changes lie outside the tested base grids. Comparisons should state which choices each condition permits. The post hoc comparison does not establish that iterative discovery was necessary (Table 1). A prospective literature-informed categorical space or an initial proposal followed by HPO could be competitive; neither was tested. Different sampled settings and training variation also prevent assigning the remaining gaps to a particular mechanism.

Acceptance and evaluation. The accepted trial records the outcome of the acceptance rule; the highest valid screening-QWK trial records the best observed score. The screening rank reversal and evaluated rejected configurations motivate reporting both trial identities and scores (Table 1; Appendix A). Applying another rule to a recorded trajectory cannot predict the changed search history; the executed ablation does not estimate policy quality across search runs (Table 2). Access restrictions, complete-essay scoring, and retained predictions support an audit, but rotating test folds still leave selection dependence unresolved.

Scoring quality. QWK improvement does not guarantee improvement on every scoring outcome (Tables 8 and 9). Independent evaluation must address the outcomes required for operational use.

6 Limitations

Study 1's loss/sampler attribution rests on its trial log, without a matched control. The studies use different agent models, and Study 2 has incomplete initialization control and an unexplained between-GPU baseline offset (§4.3.2). The extra search run repeats configuration selection, whereas nominal-seed retraining holds the selected configuration fixed; retraining does not replicate search (Bouthillier et al., 2021). The original Sonnet any-improvement run evaluates only its last accepted configuration.

Screening uses best-epoch validation scores and is optimistic by construction. Study 1 evaluates only each condition's best screening configuration, leaving the stability of its screening-to-evaluation ranking unknown. The post hoc grid, dependent test splits, scoring corrections, and missing historical sources limit claims about search superiority; they do not invalidate the reported conditional configuration comparisons.

The short-answer companion task, ASAP-SAS, was dropped after our reproduction attempt failed to match the published BERT score (Ormerod and Kwako, 2024).

7 Conclusion

In both primary comparisons, code-editing produced the configuration with the highest test-QWK point estimate. Random search over Study 2's post hoc grid recovered most of the gain over BERT (Table 1). Requiring a minimum improvement margin left accepted trials below the highest valid screening QWK in every run that accepted a trial under this rule. We recommend preserving the alternatives, acceptance decisions, and evaluation population behind each selected configuration. The primary comparisons use one search run per condition and reuse selection essays in other folds' test splits on screened prompts. Establishing search superiority requires repeated prospective searches against a literature-informed space and independent evaluation data.

References

Yoshua Bengio and Yves Grandvalet. 2004. No unbiased estimator of the variance of K-fold cross-validation. *Journal of Machine Learning Research*, 5:1089–1105.

James Bergstra, Rémi Bardenet, Yoshua Bengio, and Balázs Kégl. 2011. Algorithms for hyper-parameter optimization. In *Advances in Neural Information Processing Systems*, volume 24.

James Bergstra and Yoshua Bengio. 2012. Random search for hyper-parameter optimization. *Journal of Machine Learning Research*, 13:281–305.

Xavier Bouthillier, Pierre Delaunay, Mirko Bronzi, Assya Trofimov, Brennan Nichyporuk, Justin Szeto, Nazanin Mohammadi Sepahvand, Edward Raff, Kanika Madan, Vikram Voleti, Samira Ebrahimi Kahou, Vincent Michalski, Dmitriy Serdyuk, Tal Arbel, Chris Pal, Gael Varoquaux, and Pascal Vincent. 2021. Accounting for variance in machine learning benchmarks. In *Proceedings of Machine Learning and Systems*, volume 3, pages 747–769.

Gavin C. Cawley and Nicola L. C. Talbot. 2010. On over-fitting in model selection and subsequent selection bias in performance evaluation. *Journal of Machine Learning Research*, 11:2079–2107.

Jun Shern Chan, Neil Chowdhury, Oliver Jaffe, James Aung, Dane Sherburn, Evan Mays, Giulio Starace, Kevin Liu, Leon Maksin, Tejal Patwardhan, Lilian Weng, and Aleksander Mądry. 2025. MLE-bench: Evaluating machine learning agents on machine learning engineering. In *International Conference on Learning Representations (ICLR)*. ArXiv:2410.07095.

Jordi de la Torre, Domenec Puig, and Aida Valls. 2018. Weighted kappa loss function for multi-class classification of ordinal data in deep learning. *Pattern Recognition Letters*, 105:144–154.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.

Fabio Ferreira, Lucca Wobbe, Arjun Krishnakumar, Frank Hutter, and Arber Zela. 2026. Can LLMs beat classical hyperparameter optimization algorithms? a study on autoresearch. *Preprint*, arXiv:2603.24647.

Le Hou, Chen-Ping Yu, and Dimitris Samaras. 2017. Squared earth mover's distance-based loss for training deep neural networks. *Preprint*, arXiv:1611.05916.

Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*.

Qian Huang, Jian Vora, Percy Liang, and Jure Leskovec. 2024. MLAgentBench: Evaluating language agents on machine learning experimentation. In *International Conference on Machine Learning (ICML)*. ArXiv:2310.03302.

Kevin Jamieson and Ameet Talwalkar. 2016. Non-stochastic best arm identification and hyperparameter optimization. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, volume 51 of *Proceedings of Machine Learning Research*, pages 240–248. PMLR.

Zhengyao Jiang, Dominik Schmidt, Dhruv Srikanth, Dixing Xu, Ian Kaplan, Deniss Jacenko, and Yuxiang Wu. 2025. AIDE: AI-driven exploration in the space of code. *Preprint*, arXiv:2502.13138.

Liam Li, Kevin Jamieson, Afshin Rostamizadeh, Ekaterina Gonina, Jonathan Ben-Tzur, Moritz Hardt, Benjamin Recht, and Ameet Talwalkar. 2020. A system for massively parallel hyperparameter tuning. In *Proceedings of Machine Learning and Systems*, volume 2.

Xiaozhe Li, Jixuan Chen, Xinyu Fang, Shengyuan Ding, Haodong Duan, Qingwen Liu, and Kai Chen. 2026. OPT-BENCH: Evaluating the iterative self-optimization of LLM agents in large-scale search spaces. In *Findings of the Association for Computational Linguistics: ACL 2026*, pages 28437–28452. Association for Computational Linguistics.

Yitao Li. 2026. Greedy is a strong default: Agents as iterative optimizers. *Preprint*, arXiv:2603.27415.

Siyi Liu, Chen Gao, and Yong Li. 2025. AgentHPO: Large language model agent for hyper-parameter optimization. In *Conference on Parsimony and Learning*, volume 280 of *Proceedings of Machine Learning Research*, pages 1146–1169. PMLR.

Chris M. Ormerod and Alexander Kwako. 2024. Automated text scoring in the age of generative AI for the GPU-poor. *Chinese/English Journal of Educational Measurement and Evaluation*, 5(3). Article 5.

Christopher Ormerod and Gitit Kehat. 2025. Long context automated essay scoring with language models. In *Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con): Full Papers*, pages 35–42. National Council on Measurement in Education (NCME).

Nico Pelleriti, Sree Harsha Nelaturu, Zhanke Zhou, Zongze Li, Max Zimmer, Bo Han, and Sebastian Pokutta. 2026. What do evolutionary coding agents evolve? *Preprint*, arXiv:2605.20086.

Pedro Uria Rodriguez, Amir Jafari, and Christopher M. Ormerod. 2019. Language models and automated essay scoring. *Preprint*, arXiv:1909.09482.

Kaveh Taghipour and Hwee Tou Ng. 2016. A neural approach to automated essay scoring. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1882–1891.

Edan Toledo, Karen Hambardzumyan, Martin Josifoski, Rishi Hazra, Nicolas Baldwin, Alexis Audran-Reiss, Michael Kuchnik, Despoina Magka, Minqi

Jiang, Alisia Maria Lupidi, Andrei Lupu, Roberta Raileanu, Tatiana Shavrina, Kelvin Niu, Jean-Christophe Gagnon-Audet, Michael Shvartsman, Shagun Sodhani, Alexander H. Miller, Abhishek Charnalia, and 6 others. 2025. AI research agents for machine learning: Search, exploration, and generalization in MLE-bench. In *Advances in Neural Information Processing Systems*, volume 38.

Niki van Stein, Diederick Vermetten, and Thomas Bäck. 2025. In-the-loop hyper-parameter optimization for LLM-based automated design of heuristics. *ACM Transactions on Evolutionary Learning and Optimization*.

Benjamin Warner, Antoine Chaffin, Benjamin Clavié, Orion Weller, Oskar Hallström, Said Taghadouini, Alexis Gallagher, Raja Biswas, Faisal Ladhak, Tom Aarsen, Nathan Cooper, Griffin Adams, Jeremy Howard, and Iacopo Poli. 2025. Smarter, better, faster, longer: A modern bidirectional encoder for fast, memory efficient, and long context finetuning and inference. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2526–2547. Association for Computational Linguistics.

Orion Weller, Kathryn Ricci, Marc Marone, Antoine Chaffin, Dawn Lawrie, and Benjamin Van Durme. 2026. Seq vs seq: An open suite of paired encoders and decoders. In *International Conference on Learning Representations*.

David M. Williamson, Xiaoming Xi, and F. Jay Breyer. 2012. A framework for evaluation and use of automated scoring. *Educational Measurement: Issues and Practice*, 31(1):2–13.

Jie JW Wu, Ayanda Patrick Herlihy, Ahmad Saleem Mirza, Ali Afoud, and Fatemeh Fard. 2026. GrowthHacker: Automated off-policy evaluation optimization using code-modifying LLM agents. *ACM Transactions on Software Engineering and Methodology*. Accepted for publication.

Jiayi Xie, Kaiwei Cai, Li Kong, Junsheng Zhou, and Weiguang Qu. 2022. Automated essay scoring via pairwise contrastive regression. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 2724–2733.

Chengrun Yang, Xuezhi Wang, Yifeng Lu, Hanxiao Liu, Quoc V. Le, Denny Zhou, and Xinyun Chen. 2024. Large language models as optimizers. In *International Conference on Learning Representations (ICLR)*. ArXiv:2309.03409.

Ruosong Yang, Jiannong Cao, Zhiyuan Wen, Youzheng Wu, and Xiaodong He. 2020. Enhancing automated essay scoring performance via fine-tuning pre-trained language models with combination of regression and ranking. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1560–1569.

Helen Yannakoudakis, Ted Briscoe, and Ben Medlock. 2011. A new dataset and method for automatically grading ESOL texts. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 180–189. Association for Computational Linguistics.

Helen Yannakoudakis and Ronan Cummins. 2015. Evaluating the performance of automated text scoring systems. In *Proceedings of the Tenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 213–223.

Appendix

A Screening and selection records

The supplied halving scheduler uses 27 one-epoch starts and nine four-epoch promotions (36 evaluations, not configurations). The historical scheduler remains unauthenticated.

Table 3 records Study 1 screening; Table 4 records Study 2 comparators. The original code-editing condition's screening record is in Table 2. The grid-restricted condition's later trials used cheaper ten-epoch configurations.

Condition	Trials	Screen
code-editing (LLM)	25	**0.729**
random	27	0.681
halving	36	0.678
grid	28	0.534
TPE	27	0.692
cross-entropy start	–	0.500

Table 3: Study 1 screening record. Screen: best reduced-fidelity screening QWK; bold marks the highest. Trials include the LLM baseline and halving rung evaluations; TPE: tree-structured Parzen estimator.

Condition / configuration	Trials	Accepted	Baseline	Best	Test	SE
grid-restricted: num020 (best)	93	0	0.738	0.744	–	–
random, base grid	110	–	–	0.745	–	–
random, post hoc grid	73	–	–	0.770	–	–
Code-editing: str044 (rejected, best)	–	–	–	–	0.780	0.0030
Code-editing: str001 (rejected)	–	–	–	–	0.770	0.0030

Table 4: Study 2 screening and diagnostic evaluations. Baseline/Best: historical screening QWK. Test: corrected QWK; SE: sample fold standard error. Dashes indicate values reported elsewhere or not applicable.

The rejected diagnostic configurations str044 and str001 come from the same original code-editing search run, not independent search runs. Each test-split evaluation covers 40 prompt-fold cells. Primary test comparators are in Table 1.

B Search-run replication and retraining seeds

One additional code-editing search run uses the same instructions and starting configuration. Separately, the original accepted configuration and BERT baseline are retrained at two additional nominal seeds. The first repeats search; the second holds configuration selection fixed (Table 5).

The additional code-editing search run accepted no trial, so the original acceptance contrast was not repeated. Its best trial averages a smoothed-QWK classifier and a regression model; it narrowly missed the margin but exceeded BERT in evaluation (Table 5). Two search runs do not estimate search reliability. Separately, the original accepted configuration's evaluation gain remains positive under both nominal-seed reruns. Nominally paired seeds do not establish identical initialization (§4.3.2).

Replication diagnostic	QWK / count	SE
Search-run replicate: trials / accepted	50 / 0	–
baseline screening QWK	0.734	–
best screening QWK (r022)	0.770	–
shortfall from threshold	0.0003	–
r022 test QWK	0.771	0.0029
r022 gain over BERT	+0.019	–
BERT-base, nominal seed 7	0.755	0.0057
BERT-base, nominal seed 123	0.734	0.0149
str008, nominal seed 7	0.778	0.0035
str008, nominal seed 123	0.773	0.0024
Mean paired gain, nominal seeds 42/7/123	+0.030	–

Table 5: Search-run replication (top) and fixed-configuration retraining (bottom). Test QWK uses 40 cells per configuration; SE is sample fold standard error. Original screening records: Table 2; nominal-seed-42 test QWK: Table 1.

C Original Opus margin-rule record

Claude Opus 5 runs the code-editing instructions under the same clock and margin as Sonnet 5.

The accepted configuration is so006; the highest-screening trial is so042. Later trials exceeded the incumbent score without clearing the acceptance threshold. Recovered so042 source reads test-fold essays as unlabeled pretraining text (Appendix K); its diagnostic is reported separately and excluded from cross-condition rankings.

Opus code-editing condition	Recorded value
Trials / accepted	43 / 1
Baseline screening QWK	0.736
Last accepted / best screening QWK	0.785 / 0.794
Accepted so006 test QWK (SE)	0.788 (0.0024)
Later trials above incumbent	12
so042 transductive diagnostic QWK (SE), corrected	0.793 (0.0025)

Table 6: Original Opus margin-rule record. Test QWK and sample fold SE use 40 cells per configuration; so042 was rejected by the margin. The recovered so042 source uses test-fold text in pretraining (Appendix K).

D Input and head diagnostics

Two fixed-configuration ablations of the accepted Study 2 configuration retain only its longer input windows or its prompt-dependent head selection. The input-only configuration uses ModernBERT-base, 1600-token windows, and `CLASSIFY_MAX_CLASSES` 100 (classification on every prompt). The head-only configuration uses ModernBERT-base, 512-token windows, and `CLASSIFY_MAX_CLASSES` 10: classification for at most ten training-observed score levels and sigmoid regression otherwise. ModernBERT-512 is the common classifier comparator. Paired fold contrasts compare each component at a fixed setting of the other (Table 7); they do not allocate the BERT-relative gain into additive shares.

Head-related contrasts are larger than input-length contrasts on the common ModernBERT backbone (Table 7). The negative interaction means that the combined improvement is smaller than the sum of the two separate improvements. ModernBERT-512 itself scores below BERT. The per-prompt evidence is in Table 14.

ModernBERT configuration / contrast	QWK	SE
512-token windows, classification	0.739	0.0041
Longer windows only	0.760	0.0042
Prompt-dependent head only	0.772	0.0021
Head effect, 512-token windows	+0.032	0.005
Head effect, 1600-token windows	+0.020	0.003
Input effect, classification	+0.021	0.006
Input effect, prompt-dependent head	+0.009	0.003
Combined effect over ModernBERT-512	+0.041	–
Interaction	−0.012	0.007

Table 7: Input/head configurations and conditional QWK contrasts on ModernBERT. Combined configuration: str008 in Table 1. SE: sample fold standard error, paired by fold for contrasts.

E Secondary scoring outcomes

Condition	Exact	Adj.	MAE	SD ratio	Extr.
code-editing	0.525 (0.0064)	0.826 (0.0018)	0.94 (0.0098)	0.951 (0.0141)	0.437 (0.0295)
random	0.532 (0.0027)	0.817 (0.0014)	1.02 (0.0117)	0.926 (0.0086)	0.457 (0.0150)
halving	0.528 (0.0043)	0.815 (0.0040)	1.02 (0.0130)	0.905 (0.0109)	0.399 (0.0126)
grid	0.544 (0.0037)	0.823 (0.0018)	1.02 (0.0173)	0.872 (0.0037)	0.403 (0.0112)
TPE	0.515 (0.0074)	0.795 (0.0053)	1.12 (0.0561)	0.813 (0.0351)	0.366 (0.0280)

Table 8: Study 1 secondary outcomes: means over 40 cells, with sample fold SE in parentheses. Adj.: adjacent agreement; MAE: mean absolute error in score points; SD: standard deviation; Extr.: extreme-score agreement. Definitions: §3.4; QWK: Table 1.

Configuration	Exact	Adj.	MAE	SMD	SD ratio	Extr.
BERT-base (baseline)	0.529	0.830	0.948	+0.013	0.921	0.388
ModernBERT-512 (baseline)	0.521	0.814	1.037	+0.024	0.938	0.448
str008 (accepted)	0.527	0.836	0.868	+0.033	0.933	0.409
str044 (best)	0.525	0.836	0.876	+0.031	0.945	0.433
str001 (rejected)	0.520	0.836	0.889	+0.017	0.918	0.374
num020 (grid-restricted)	0.531	0.825	0.952	+0.054	0.907	0.410
random, base grid	0.529	0.825	0.964	+0.037	0.915	0.400
random, post hoc grid	0.524	0.832	0.881	+0.021	0.936	0.393
r022 (replicate, best trial)	0.511	0.832	0.917	+0.061	0.960	0.389
k042 (any-improvement)	0.508	0.827	0.927	−0.011	0.977	0.437
Fold SE range	0.002–0.009	0.001–0.005	0.005–0.028	–	0.004–0.019	0.005–0.024

Table 9: Legacy Study 2 secondary outcomes: means over 40 prompt-fold cells, nominal seed 42. Adj.: adjacent agreement; MAE: mean absolute error in score points; SMD: standardized mean difference; SD: standard deviation; Extr.: extreme-score agreement (§3.4). SE: sample fold standard error; dash: unreported. SE ranges retain the original full cohort, including omitted configurations. Exact-agreement contrast against BERT: str008 −0.002 (paired SE unreported); k042 −0.021 (paired fold SE 0.005).

F Hyperparameter grids

Per-axis values for the fixed grids of §3.2.1 and §3.2.2.

Study 1, 864 points ($4 \times 3 \times 3 \times 3 \times 2 \times 2 \times 2$): learning rate $\{5 \times 10^{-5}, 1 \times 10^{-4}, 2 \times 10^{-4}, 3 \times 10^{-4}\}$; LoRA rank $\{8, 16, 32\}$; LoRA alpha $\{8, 16, 32\}$; dropout $\{0.0, 0.1, 0.2\}$; warmup steps $\{0, 100\}$; scheduler {linear, cosine}; weight decay $\{0.0, 0.01\}$.

Study 2, 192 points ($4\times3\times2\times2\times2\times2$): learning rate $\{1\times10^{-5}, 2\times10^{-5}, 3\times10^{-5}, 5\times10^{-5}\}$; batch size $\{9, 16, 32\}$; epochs $\{10, 20\}$; warmup ratio $\{0.0, 0.1\}$; weight decay $\{0.0, 0.01\}$; label smoothing $\{0.0, 0.1\}$.

Study 2 post hoc grid, 192 points ($2 \times 2 \times 3 \times 2 \times 2 \times 2 \times 2$): backbone {BERT-base, ModernBERT-base}; max length $\{512, 1600\}$; head rule {regression everywhere, classification up to 10 score levels, classification everywhere}; learning rate $\{2 \times 10^{-5}, 5 \times 10^{-5}\}$; epochs $\{10, 20\}$; weight decay $\{0.0, 0.01\}$; label smoothing $\{0.0, 0.1\}$. The backbone, input-length, and head-rule choices come from the accepted configuration and were added to the grid after the code-editing condition finished; batch size and warmup ratio are pinned at the starting configuration's values. Screening counts are in Table 4; the logs omit the head threshold, so distinct full configurations cannot be counted from them. This post hoc revised grid is not a superset of the original numeric grid.

G Minimum improvement margin, acceptance thresholds and standard errors

Table 10 records the uncertainty estimates used to motivate the margins and the historical screening values used to inspect rejected improvements. Cell-based SEs pool heterogeneous prompt-fold cells, not independent search runs; Table 1 retains the distinct logged propagated SE.

Minimum improvement margins. Study 1's 0.02 rule used a baseline estimate from five training-seed reruns of the same configuration. Study 2 used five random seeds of the unedited baseline at the code-editing-condition setting; its screening spread was dominated by prompt 8. The required improvement is $2\sigma\sqrt{2} = 0.036$. The cross-GPU baseline offset (Table 11) is consistent with, but does not validate, that width.

Incumbent scores and rejected improvements. Table 10 records the renewed thresholds after str008 and so006 and the later above-incumbent trials. In the separately executed acceptance-rule condition, k007 falls below the threshold the original margin would have imposed.

Record	Configuration / statistic	Value
Study 1 cell-based SE	code-editing / random	0.0104 / 0.0122
	halving / grid / TPE	0.0133 / 0.0222 / 0.0379
Study 1 sample-SD propagation	code-editing / random	0.0052 / 0.0046
Study 1 baseline seed estimate	σ	≈ 0.0073
Study 2 baseline screening QWK	range / mean	0.714–0.740 / 0.729
Baseline screening seed SD	overall / prompt 8	0.0128 / 0.031
Code-editing incumbent	str008	0.7656
Renewed acceptance threshold	after str008	0.8016
Later above-incumbent trials	count	4
	str012 / str034	0.7663 / 0.7658
	str035 / str044	0.7659 / 0.7689
Opus incumbent / threshold	after so006	0.785 / 0.821
Any-improvement trial / counterfactual threshold	k007	0.769 / 0.777
str008 screening QWK	seeds 42 / 7	0.766 / 0.750
	seeds 123 / 2026	0.755 / 0.753
	seed mean	0.756

Table 10: Historical margin diagnostics. Screening values are QWK; SD: standard deviation; SE: standard error. Estimator definitions follow.

Standard errors. Study 1 rows in Table 1 retain the logged convention: per-prompt SE = population SD of 5 fold QWKs$/\sqrt{5}$; 8-prompt SE = root sum of squares over prompts$/8$. Study 2 tables: sample SD of the 5 fold means$/\sqrt{5}$. The cell-based SEs in Table 10 use sample SD over 40 cells$/\sqrt{40}$. Paired Study 2 contrasts use sample SD of five fold-wise differences$/\sqrt{5}$. These describe the observed configurations; fold dependence and variation across search runs are not modeled. Table 10 also gives the Study 1 LLM and random estimates under sample-SD propagation rather than the logged convention.

All-essay QWK correction (Study 2). For integer truth scores y and predictions $\hat{y}$, we use the complete contiguous score scale, equivalently

$$1 - \frac{\overline{(y - \hat{y})^2}}{\mathrm{Var}(y) + \mathrm{Var}(\hat{y}) + (\bar{y} - \bar{\hat{y}})^2},$$

with population moments over every essay in the prediction file. Historical records are retained unchanged. Audit counts and the independent QWK cross-check are in Table 11 and §4.3.2. Corrected scoring evaluates the already selected predictions; it does not redo historical validation or search.

The recovered legacy analyses use the QWK and fold SE in `study2_metrics.py`.

Audit record	Value
Legacy evaluation entries / valid evaluation cells	20 / 800
Cells with predictions / discrepant QWK	800 / 60
Original audit cells / discrepant QWK	680 / 51
Omitted extreme essay: prompt / fold / score	7 / 0 / 2
	8 / 1 / 10
	8 / 4 / 60
Published BERT comparison QWK	0.7475
Historical logged BERT mean	0.752
Between-GPU baseline screening offset	0.014

Table 11: Legacy evaluation records and all-essay audit, including recovered evaluations. Published BERT comparison: Rodriguez et al. (2019). Corrected BERT mean: Table 1; corrected so042 outcome: Table 6.

H Study 2 agent instructions (`program.md`), verbatim

Karpathy's loop has three files: `prepare.py` (protected data and evaluation code), `train.py` (the one file the agent edits), and `program.md` (agent instructions). The loop minimizes validation bits per byte, with a five-minute training limit per trial; any improvement advances a git branch.

We reproduce the operating prompt read by the code-editing condition in Study 2 (§3.2.2) verbatim except for re-wrapping long lines, transliterating Unicode symbols (arrows, $\leq$, $\neq$, σ) to ASCII for the monospace font, and renaming the experimental-branch term and its command-line flag to *condition* and `-condition`, for consistency with this paper's terminology. These are editorial changes, not the original executable interface. The excerpt retains the historical labels "structural", "numerical", "keep", and "bar"; the paper calls these code-editing, grid-restricted, accepted, and incumbent score, respectively. Edit permissions and margins are specified in §3.2.1 and §3.2.2. Study 1's supplied instruction variants also differ in paths, training instructions and confirmation-rerun permissions; exact historical variants remain unauthenticated.

```
# autoresearch loop -- ASAP-AES from a strong baseline (LLM
  condition, structural)

Goal: get the highest `pooled_val_qwk` (mean val-split QWK over
  ASAP-AES sets 2,3,5,8, computed by frozen `prepare.py`) within 12
  h wallclock.

## Setup (once)
- You are in `/mnt/data/jufrost/loop_bench_t2/structural/` (its own
  git repo). `python` and `trackio` on PATH are the orchestrator
  venv's; the GPU number comes from the launch prompt. Trial ids are
  `str000`, `str001`, ... Do not spawn subagents; work in this
  session only.
- Read: `train.py` (you edit this -- the whole file except
  `main()`), `prepare.py` (frozen: data, splits, seed, metric),
  `run.py` (frozen driver; `SPACE` lists the numeric knobs the
  classical condition searches).
- `python run.py --condition autoresearch --start` (stamp the 12 h
  clock).
- Baseline = trial `str000`: `python run.py --condition autoresearch
```

  --gpu <GPU> --trial-id str000` with `train.py` unedited; its
  `pooled_val_qwk` is the starting bar. It is the reproduced
  Rodriguez-Jafari-Ormerod 2019 BERT-base recipe and matches the
  published bar (8-set mean ~0.75) -- you are improving the
  literature baseline, not a stock model. The backbone (`MODEL`) is
  yours to change.
- `git branch autoresearch/keep-decisions` at this tree.

Experimentation
Run one trial: `python run.py --condition
  autoresearch --gpu <GPU> --trial-id str<NNN> [--config '<json>']`
  (~15 min: 4 sets x 20 epochs; a cheaper recipe runs faster).

Rule	Value
CAN edit	anything in `train.py` above `main()`: model/head,
loss, optimizer, schedule, batch size, epochs, windowing/decoding,
data handling, regularization. `--config` may also set the
constants listed in `HP_NAMES`.
CANNOT edit	`train.py`'s `main()` (result assembly),
`prepare.py`, `run.py`, the seed, the splits, the metric,
dependencies (no new packages). The test split is unreadable and
off-limits.
Metric	`pooled_val_qwk` (val split, best-epoch selection on
val), higher better
Keep margin	**+/-0.036** (measured: sigma of the unedited
baseline over 5 seeds at this exact setting) -- a keep must beat
the bar by MORE than this; within it = no change = discard. No
reruns: the seed is fixed, a rerun is a replicate.
Budget	12 h wallclock incl. your thinking time; hard stop at
clock expiry. A cell over `prepare.CELL_HARD_CAP_S` exits 124.
Tie-break	simpler wins; a win from deleting code is best

Starting point
The baseline is the published recipe (CLS
  classification over the set's discrete scores, cross-entropy,
  AdamW 5e-5, batch 9, 20 epochs, 510-token overlapping windows,
  best-on-val epoch). There are **no ASAP-specific priors** beyond
  that -- discover what moves QWK on this data from scratch. QWK is
  quadratic-weighted: far misses cost more than near ones.

Outcome -- read by EXIT CODE, then Trackio (never parse stdout)
Exit	Meaning	Action
0	trial done	`trackio get run --project loop_bench_t2 --run
autoresearch_str<NNN> --json` -> `pooled_val_qwk`, `valid`,
per-set `qwk_s<k>`; cross-run compare with `trackio list runs
--project loop_bench_t2`. `valid=0` (no metric key) -> invalid,
discard.
124	clock expired OR a cell hit the hard cap	run `python
run.py --condition autoresearch --check`: exit 1 -> clock done,
STOP; exit 0 -> the trial was too slow, discard
2	no clock	you skipped `--start`; run it once

```
| other | crash | read `runs/autoresearch_str<NNN>/train.log`; fix
  if trivial and re-run the same commit <=2x, then move on |

## The loop -- RUN UNTIL THE CLOCK EXPIRES
1. `python run.py --condition autoresearch --check` -> exit 1
  (expired) -> STOP. `remaining_h` < 0.4 -> STOP (a trial would not
  finish).
2. Bar = best kept `pooled_val_qwk`. Commit or revert stray
  *tracked* changes (never `git clean`; `runs/` and
  `clock_start.json` are gitignored and must survive), then `git
  checkout -b autoresearch/experiment/str<NNN>
  autoresearch/keep-decisions`.
3. Form ONE idea (a diagnosis of the last result, a combination of
  keeps, an ablation). Prefer moves the numeric grid in `run.py`
  cannot express.
4. Edit `train.py`; `git commit` (message = the one-line why).
5. Run the trial; read the outcome by exit code (table above).
6. Admissibility first: exit != 0 -> handle per table. A keep is
  `valid` AND beats the bar by **more than the margin**.
7. Keep -> `git branch -f autoresearch/keep-decisions HEAD`. Discard
  -> leave the branch; the next idea branches off the bar again.
  Nothing is lost (every trial is in Trackio).

**NEVER STOP** before the clock. Out of ideas -> think harder: a new
  loss family, a different head or pooling, a different decode of
  window predictions, a sampler, a data transform, a schedule. The
  human is asleep; keep finding wins over a strong baseline.
```

I Study 1 trial log

Table 12 selects baseline 000, the accepted trials (001, 002, 004), and the highest-screening trial (005) from the Study 1 LLM trial log (mean 4-prompt screening QWK). Historical trial descriptions and outcomes retain their original vocabulary, including "keep" and "transferred". Trial 005 exceeded the last accepted trial 004 but missed the margin; it is the configuration evaluated in Table 1.

#	idea	QWK	outcome
000	baseline cross-entropy	0.500	(base)
001	soft-ordinal CE (τ=1)	0.640	**keep** +0.14
002	squared EMD (param-free)	0.673	**keep** +0.033
004	EMD + sqrt-inv-freq sampler	0.715	**keep** +0.042
005	004 + lr= 2×10^{-4}	0.729	sub-margin; transferred

Table 12: Selected Study 1 trial records. QWK: mean screening quadratic-weighted kappa; outcome text is historical. Notation key and authentication limits follow the table.

Historical notation key. CE: cross-entropy; EMD: Earth Mover's Distance; lr: learning rate; τ: the recorded soft-ordinal loss parameter. The particular τ construction remains unauthenticated.

J Per-prompt test-split QWK

The historical code-editing screening rows in Table 14 show the largest gain on prompt 8.

Condition	P1	P2	P3	P4	P5	P6	P7	P8
code-editing	0.807	0.668	0.704	0.826	0.802	0.825	0.808	0.726
random	0.792	0.674	0.717	0.821	0.814	0.819	0.766	0.615
halving	0.773	0.658	0.700	0.826	0.794	0.766	0.779	0.613
grid	0.753	0.655	0.711	0.808	0.796	0.794	0.806	0.396
TPE	0.402	0.512	0.682	0.802	0.805	0.811	0.600	0.617

Table 13: Study 1 per-prompt test QWK, averaged over five folds. P1–P8: essay prompts.

Config.	P1	P2	P3	P4	P5	P6	P7	P8	P8 SE
BERT-base	0.798	0.653	0.688	0.825	0.807	0.809	0.811	0.629	0.0292
ModernBERT-512	0.782	0.672	0.689	0.804	0.819	0.807	0.790	0.553	0.0228
str008 (accepted)	0.832	0.686	0.699	0.817	0.816	0.813	0.834	0.749	0.0146
str044 (best)	0.830	0.706	0.678	0.812	0.821	0.810	0.839	0.743	0.0133
str001 (rejected)	0.810	0.683	0.685	0.807	0.810	0.806	0.836	0.723	0.0161
num020 (grid-restricted)	0.791	0.645	0.686	0.820	0.811	0.804	0.814	0.624	0.0361
random, base grid	0.794	0.674	0.677	0.817	0.804	0.794	0.816	0.630	0.0234
random, post hoc grid	0.827	0.677	0.682	0.816	0.812	0.808	0.835	0.749	0.0249
r022 (replicate)	0.814	0.662	0.697	0.817	0.807	0.813	0.832	0.729	0.0139
k042 (any-impr.)	0.828	0.700	0.665	0.813	0.807	0.816	0.828	0.748	0.0191
longer windows only	0.819	0.678	0.696	0.809	0.806	0.806	0.779	0.690	0.0193
head only	0.810	0.686	0.683	0.815	0.807	0.815	0.830	0.727	0.0197
Historical Sonnet code-editing screening QWK (screened prompts only)									
Baseline	–	0.739	0.716	–	0.817	–	–	0.622	–
str008	–	0.752	0.724	–	0.822	–	–	0.765	–

Table 14: Legacy Study 2 per-prompt test QWK (five-fold mean) and historical screening QWK; seed 42 unless stated. P1–P8: prompts; P8 SE: sample standard error over five prompt-8 folds. Any-impr.: any-improvement. Dashes indicate unreported values.

K Evaluated configurations

Historical essay-to-split assignments, search access to test data, when evaluation results informed later conditions, and complete executed-code identities remain unauthenticated. Recovered selected sources narrow the artifact gaps but do not authenticate every historical execution. The post hoc head threshold and the complete evaluator environment also remain unverified. A recovered starting script matches the code-editing launch hash but creates the model before applying its explicit training seed; nominal seed pairing therefore does not establish identical initialization.

Our implementation uses `bert-base-uncased`, a classification head on the encoder's summary (CLS) token over each prompt's discrete scores, cross-entropy, AdamW at 5×10^{-5}, a batch size of 9, 20 epochs, and 510-token overlapping windows, with the epoch selected by validation QWK.

Tables 15 and 16 list the configurations evaluated in Results, with logged hyperparameters and reported commit identifiers.

Condition	LR	Rank	Alpha	Drop.	Warmup	Sched.	WD
code-editing	2e-4	16	16	0.2	100	linear	0.0
grid	5e-5	8	16	0.0	0	linear	0.01
random	2e-4	16	32	0.1	0	linear	0.01
TPE	3e-4	8	32	0.0	0	cosine	0.01
halving	3e-4	8	16	0.1	100	linear	0.0

Table 15: Study 1 evaluated configurations (Table 1). LR: learning rate; Drop.: dropout; Sched.: scheduler; WD: weight decay; Rank/Alpha: LoRA settings. Warmup is in steps.

All four Study 1 HPO conditions train with cross-entropy. The code-editing condition set only its learning rate; its other LoRA values are unedited harness defaults. Its code change (commit `89c50543`) is the EMD loss plus sampler of §4.1.1 (Table 15).

Config.	Backbone	Input	Head	Commit
BERT-base	BERT	512, wind.	class.	–
ModernBERT-512	ModernBERT	512, wind.	class.	–
str008 (accepted)	ModernBERT	1600, wind.	prompt-dep.	`3a1493b`
str044 (best)	ModernBERT	1600, wind.	prompt-dep.	`8ee6745`
str001 (rejected)	BERT	512, wind.	regr.	`661af18`
num020 (grid-restricted)	BERT	512, wind.	class.	–
random, base grid	BERT	512, wind.	class.	–
random, post hoc grid	ModernBERT	1600, wind.	unverified	–
r022 (replicate)	BERT	512, wind.	ensemble	`aebdf00`
k042 (any-impr.)	ModernBERT	1536, wind.	regr.+rank	`e48fe43`
so006 (Opus 5)	ModernBERT	1536, wind.	regr.+quant.	`b54ab82`
so042 (transductive)	ModernBERT	1536, wind.	aux.+quant.	unverified
longer windows only	ModernBERT	1600, wind.	class.	–
head only	ModernBERT	512, wind.	prompt-dep.	–

Table 16: Study 2 recorded configurations. Wind.: overlapping windows; class.: classification; regr.: regression; prompt-dep.: prompt-dependent head selection; rank: pairwise ranking loss; quant.: quantile score discretization; aux.: auxiliary trait heads; any-impr.: any-improvement.

All configurations use learning rate 5e-5, 20 epochs, and batch size 9, except k042 (learning rate 3e-5) and random, post hoc grid (10 epochs). Recovered k042, so006, and so042 sources request bfloat16 (BF16); historical executed precision is unauthenticated.

Window lengths include two special tokens. Recovered sources split overlength essays into windows, with the final window overlapping; essays that fit use one window. Heads: "class." is the source classification head over each prompt's discrete scores; "regr." is sigmoid-regression with mean squared error (MSE); "prompt-dep." uses classification for prompts with ≤ 10 training-observed score levels and regression otherwise. "regr.+rank" uses weighted scalar Smooth L1 loss ($\beta = 0.1$) with pairwise ranking loss (weight 0.5, margin 0.05) in the recovered k042 source.

The recovered so006 and so042 sources use quantile-based score discretization: they rank the evaluated essays' continuous scores and allocate consecutive ranks to classes using rounded cumulative training-score proportions multiplied by the evaluated essay count. The allocation therefore depends on the evaluated collection of essays. Stable sorting retains the incoming order of tied scores; ties can span class boundaries. The historical ordering of essays remains unauthenticated.

The classifier/regression ensemble description applies to r022. Its recovered source converts each window's classifier argmax to a score and clips the raw regression output to the class-index range before interpolating between training-observed score values. It averages each member's scores over an essay's windows, then equally averages the two per-essay values on the human-score scale before rounding to a score class. This regression conversion uses no sigmoid. The rounding helper's tie behavior remains unverified.

str044 adds classifier dropout 0.1 and weight decay 0.02 to str008. num020 adds weight decay 0.01 and label smoothing 0.1. Random search on the base grid is trial 86 and random search on the post hoc grid is trial ID 29 (condition counts in Table 4). The post hoc random settings omit the head threshold; the earlier description as regression everywhere cannot be authenticated.

The recovered so042 source adds auxiliary trait heads and masked-language-model pretraining; it defaults to three pretraining epochs, whereas the screening commit describes six and the evaluation manifest supplies no override. Its pretraining always reads fold-0 training text across all prompts. Comparison with the prediction IDs shows that this pool contains every test essay in folds 1–3, and none in folds 0 or 4. This is unlabeled test-text exposure, not supervised use of test labels. The reported so042 aggregate therefore describes a transductive evaluation with unresolved screening-source identity; it is reported separately and excluded from cross-condition rankings.

LLM Difficulty Prediction Across Three L2 Task Types and Two Examination Forms

Meng Lyu Boston College
Moses Oluoke Omopekunola HSE University
Feng Wang University of Massachusetts Amherst

Abstract

Large language models (LLMs) have been proposed as a substitute for field trials in item-difficulty estimation, but evidence about where they succeed is usually drawn from a single test form. We compare LLM difficulty predictions across three task types that make different language demands, grammar cloze, semantic cloze and reading comprehension, on two operational Grade 9 English examinations sat by the same cohort in China (1,199 and 1,166 examinees). Three LLMs estimated difficulty for 60 MCQ items under the same standardized protocol. On Form 1, grammar cloze had the highest Pearson correlation for all three models; the Fisher-averaged correlations were .81 for grammar cloze and .34 for semantic cloze. On Form 2, reading had the highest correlation for all three models (Fisher-averaged .78), and the grammar–semantic difference was smaller and inconsistent in direction. With ten items per task type, a task-type advantage seen on one form needs testing on other forms before it is read as a property of the model.

1 Introduction

Knowing how difficult an item is before administration supports test assembly, item-bank triage and adaptive testing, yet empirical difficulty is normally recovered only after costly field trials (Benedetto et al., 2023). Traditional readability formulas and NLP indices predict difficulty weakly, especially for second language (L2) items, where difficulty depends on construct-specific demands that surface features miss (Benedetto et al., 2023; Hwang et al., 2025).

Large language models offer an alternative, and several studies report moderate correlations between LLM difficulty estimates and empirical difficulty across domains and prompting strategies (Caines et al., 2023; Raina and Gales, 2024; Razavi and Powers, 2025; Kolesnikova et al., 2026). Beyond overall prediction accuracy, a natural next question concerns how model performance varies across different types of items.

We compare LLM difficulty predictions across grammar cloze, semantic cloze and reading comprehension, and examine whether task-type differences persist across two examination forms. The three task types are chosen because they make different language demands. A grammar cloze blank is resolved by applying a morphosyntactic rule; a semantic cloze blank requires recovering meaning from the surrounding context; a reading item requires integrating information across sentences and, often, inference beyond the literal text. The two cloze sections carry an additional advantage for comparison: they share a four-option format and appear in the same paper, so a difference between them is not a difference of response format.

Evidence on this question is usually produced from one test form, a few dozen items, one prompt and one run per item, with sub-group claims read off subsets of ten or twenty items. We therefore ask the same question on two forms of the same testing program, with three models and three runs per item, under one documented protocol.

This paper contributes: (i) a two-form, three-model, three-run comparison of LLM difficulty validity across three L2 task types on an operational Chinese Grade 9 English examination; (ii) a psychometric audit of a four-administration item bank documenting when real difficulty can serve as a criterion; and (iii) a documented protocol, prompt set and per-item prediction record.

2 Aims

This study examines whether LLM predictions of item difficulty vary across grammar cloze, semantic cloze, and reading comprehension, and whether these patterns are consistent across two examination forms from the same testing program. We

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 441–447
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

consider both relative validity, defined as how well predictions track differences in empirical difficulty, and absolute accuracy, defined as how closely predicted $p$roportions correct match observed values. The study addresses two research questions:

RQ1. How well do three LLMs predict empirical difficulty in each of the three task types?

RQ2. Do the differences between task types observed on one test form also appear on a second form from the same testing program?

3 Data

3.1 Two examination forms

The primary data are two operational English examinations sat by a single Grade 9 cohort (approximately CEFR B1) at a public school in a western city in China. Form 1 is the mid-term; Form 2 is a subsequent mock examination. From each form we analyze 30 four-option items: grammar cloze, semantic cloze and reading comprehension. Form 1 was taken by 1,199 students, of whom 1,168 answered every analyzed item; item difficulties are computed on all valid responses to each item (n = 1,192 to 1,199). For Form 2, after removing rows with no responses to the analyzed items and one class-average summary row, 1,166 examinees answered every analyzed item, and every item difficulty is based on those 1,166 responses.

Each form supplies one passage for the grammar cloze section, one for the semantic cloze section, and three for the reading section. Task type and passage are therefore confounded within a form. Using a second form gives each task type a second set of passages, which reduces the chance that a pattern depends on one particular text, but does not separate task type from passage (§7).

The two forms were written to different difficulty targets. Form 1 is a routine mid-term, set at the difficulty of everyday classroom testing; Forms 2 to 4 are mock papers set at the difficulty of the high-school entrance examination itself. This follows the regional norm in the western city in China, where teaching materials and routine assessment are deliberately pitched above the entrance examination, so that candidates meet an easier paper on the day.

3.2 Difficulty range across the four administrations

The same cohort sat four curriculum-sequential forms. Table 1 summarises all four and Figure 1 shows the two that enter the analyses, because the difficulty range of a form determines whether it can serve as a criterion at all. After excluding rows with no responses to the analysed items, and one class-average summary row present in each of the Form 2–4 exports, the numbers of examinees answering every analysed item are 1,168, 1,166, 1,162 and 1,151.

Form 1 spans the full difficulty range. Forms 2–4 are compressed against the ceiling, with 60–76% of items above p = .90 and 34–47% of examinees answering every analysed item correctly. Mean $SE(p)$ in those forms is .008, so item difficulty is estimated precisely over a narrow range. KR-20 for the analysed items remains high in all four forms (.87–.90), because the full-sample total-score variance is large, while the marginal reliability of a unidimensional 2PL model fitted to the same items by marginal maximum likelihood falls from .85 in Form 1 to .65–.68 in Forms 2–4. The two indices answer different questions and we report both.

We use Forms 1 and 2 for the LLM analyses and treat Forms 3 and 4 as documentation of the difficulty range. Forms 3 and 4 replace the two ten-item cloze sections with a single fifteen-item cloze, so they do not support a grammar–semantic comparison within the same paper.

4 Method

4.1 Criterion and scale

The criterion is the classical proportion correct p for each item in its own form. The primary scale is proportion correct; correlations, bias and absolute error are reported on it. As a secondary check we repeat the correlations on the logit scale, transforming predicted and empirical values in the same direction; predictions are bounded to $[.01, .99]$ by the prompt and the largest empirical p is .989, so no boundary handling is required.

Because the two forms were written to different difficulty targets and share no items, we do not place their difficulty parameters on a common scale and we do not pool raw difficulty. All primary analyses are within form. A secondary pooled analysis standardises empirical and predicted $p$roportion correct within each form-by-

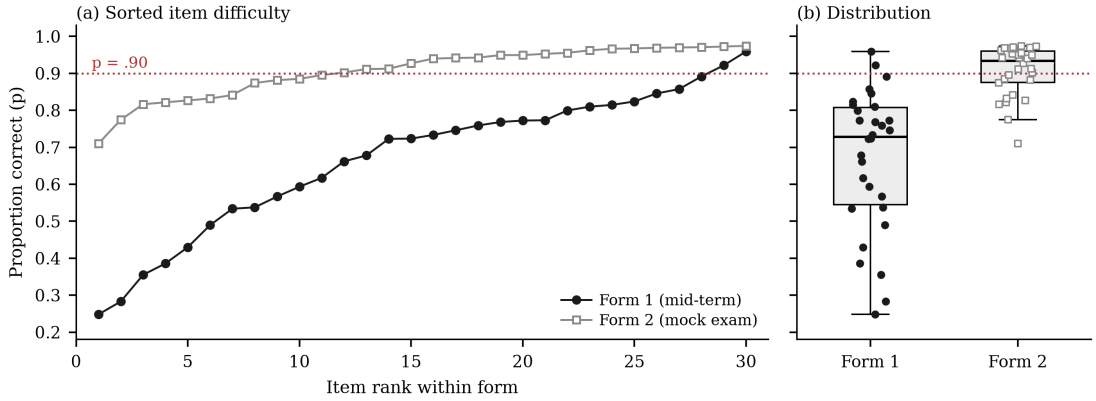


Figure 1: Empirical item difficulty across the two examination forms. Panel (a) plots proportion correct sorted within each form; panel (b) shows the corresponding distributions with every item overlaid. Only Forms 1 and 2 enter the LLM analyses.

Form	Analysed	Cloze	Reading	n complete	Mean p (SD)	Range of p	KR-20
Form 1 (mid-term)	30	20	10	1,168	.67 (.19)	.25–.96	.87
Form 2 (mock)	30	20	10	1,166	.91 (.07)	.71–.97	.90
Form 3	25	15	10	1,162	.89 (.08)	.65–.98	.89
Form 4	25	15	10	1,151	.92 (.07)	.72–.99	.87

Table 1: Item bank composition and empirical difficulty by administration. p = classical proportion correct. n complete = examinees answering every analysed item; item p is computed on all valid responses to each item, and KR-20 on complete responses. Forms 1 and 2 supply the LLM analyses; Forms 3 and 4 document the difficulty range only.

task-type cell before pooling (§5.2).

4.2 Prompting protocol

The protocol was written and fixed before any model was called. Each item is one independent call; no conversation state is carried between items. The prompt contains the full passage, the sentence containing the target blank or the reading stem, and the four options, together with a fixed context block naming the population (Grade 9 Chinese EFL learners, approximately CEFR B1), the task (estimate the expected proportion correct), and the required JSON response. All cloze blanks in the passage are shown open, as the students saw them. Chinese glosses printed in the original papers are removed from the model input, so that the two forms are presented on the same terms.

The same template is used for every item and every model, so the only thing that varies between calls is the item. The model is not told which option is correct, so the estimate is not conditioned on the answer. The prompt carries no expert metadata such as the grammar point, its position in the national curriculum, or the position of the blank in the passage: such features exist for some task types and not others, and supplying them would make the task types differ in the information given rather than in the item itself. The prompt also does not name the task type, since that label names one of the comparisons under test.

Three models were used: Claude Opus 4.8, GPT-5.5 and Gemini 3.1 Pro. Each item was run three times per model at default sampling settings; the reported prediction is the mean of the three runs and we report the within-item standard deviation across runs. Raw responses, resolved version strings and parameters are archived for all 540 calls.

Sixty items were run. Form 2 item 50 asks for the order of an experiment whose four options are sequences of three pictures labelled a, b and c; the pictures were not supplied to the models, which received only the letter sequences, so the item is excluded from all analyses. The main analyses therefore use 59 items.

5 Results

5.1 Validity by task type and form

Table 2 and Figure 2 give the main result. On Form 1, all three models produced their highest

Pearson correlation for grammar cloze (r = .72, .83, .85), and for every model grammar cloze exceeded semantic cloze. The models did not agree beyond that: two ranked reading above semantic cloze and one ranked semantic cloze above reading. Averaged across models by Fisher z, the Form 1 values are grammar .81, reading .46 and semantic .34.

On Form 2, all three models produced their highest Pearson correlation for reading (r = .80, .83, .71). The grammar–semantic difference was smaller than on Form 1 and not consistent in direction: two models placed grammar above semantic cloze and one placed semantic cloze above grammar. The Fisher means are reading .78, grammar .37 and semantic .23.

Part of the change in the cloze sections is consistent with restriction of range: the standard deviation of empirical p falls between the two forms from .257 to .064 for grammar cloze and from .170 to .044 for semantic cloze. For reading the standard deviation also falls (from .142 to .089) while the coefficients rise, so range alone does not account for the pattern. Every confidence interval in Table 2 is wide, and most of them overlap.

5.2 The grammar–semantic contrast across forms

Table 3 expresses the grammar-versus-semantic comparison as a Fisher z difference within each form and model. On Form 1 the difference is positive for all three models (Δz = 0.89, 0.53, 0.86). On Form 2 it is smaller and, for one model, negative (0.35, −0.08, 0.19). Every interval in the table includes zero.

The second form therefore gives limited support to the Form 1 pattern. We do not treat the three models as independent samples: they predict the same items, so their errors are shared. Two sets of passages per task type cannot support a passage-level inference.

Pooling the two forms, after standardising empirical and predicted $p$roportion correct within each form-by-task-type cell, gives 20 items per cloze task type and 19 for reading. The pooled correlations are .56 to .62 for grammar cloze, .47 to .74 for reading and .05 to .47 for semantic cloze across the three models: the separation between grammar cloze and reading that is visible within each form does not survive pooling. We report this as a secondary summary, since two sets of passages per task type do not constitute passage-level replication.

5.3 Calibration

Relative validity and absolute calibration are separate questions, and both matter for triage. Table 4 reports calibration by form and model, and Figure 3 plots every analysed item.

On Form 1, where empirical difficulty varies widely, all three models compress the range: predicted difficulty carries 40 to 53 per cent of the dispersion of the empirical values, and two of the three place items too easy on average. On Form 2, where empirical difficulty is itself narrow, dispersion and bias vary by model.

The last column of Table 4 gives the error of a constant prediction at the form mean. Only one model is more accurate than that baseline on each form. Because the baseline uses labels from the same sample it is a descriptive reference rather than a predictor anyone could deploy, but it shows that a correlation of .7 or .8 with empirical difficulty need not yield usable absolute accuracy.

6 Discussion

Across the two forms, the task type for which LLM difficulty estimates correlated most strongly with empirical difficulty was not the same. On the first form all three models produced their highest correlation for grammar cloze, the section whose difficulty is governed by a morphosyntactic rule visible in the item text, and each of them produced a lower correlation for the meaning-dependent cloze section. That pattern is compatible with an account in terms of learner interlanguage: meaning-based difficulty may depend on which lexical and inferential associations a particular population has consolidated, which the item text does not carry. On the second form all three models produced their highest correlation for reading, and the grammar–semantic difference was smaller and inconsistent in direction.

We are not in a position to say which picture is the more representative. Each rests on ten items per task type, and the intervals in Table 2 are wide and mostly overlapping. The reading we would offer is narrow: a task-type advantage observed on a single form is a description of that form, and needs testing on further forms before it is used to decide where automated difficulty triage can be trusted.

Reading deserves a separate note. On Form 2 all three models produced their highest correlation

Form	Model	Grammar cloze	Semantic cloze	Reading	All items	Mean run SD
Form 1	Claude Opus 4.8	.72* [.16, .93]	.02 [−.62, .64]	.32 [−.39, .79]	.39	.017
	GPT-5.5	.83** [.42, .96]	.58 [−.07, .89]	.15 [−.53, .71]	.44	.022
	Gemini 3.1 Pro	.85** [.46, .96]	.37 [−.34, .81]	.77** [.26, .94]	.64	.022
	Mean (Fisher z)	*.81*	*.34*	*.46*	*.50*	
Form 2	Claude Opus 4.8	.41 [−.30, .82]	.08 [−.58, .68]	.80* [.28, .96]	.44	.014
	GPT-5.5	.29 [−.41, .78]	.36 [−.35, .81]	.83** [.36, .96]	.52	.017
	Gemini 3.1 Pro	.40 [−.31, .82]	.23 [−.46, .75]	.71* [.08, .93]	.54	.020
	Mean (Fisher z)	*.37*	*.23*	*.78*	*.50*	

Table 2: LLM difficulty validity by task type and form. Pearson correlations between predicted and empirical *p*roportion correct, 95% CIs in brackets. k = 10 per cell, except Form 2 reading (k = 9, picture item excluded). $^{*}p < .05$, $^{**}p < .01$. Run SD is the mean within-item standard deviation across three runs.

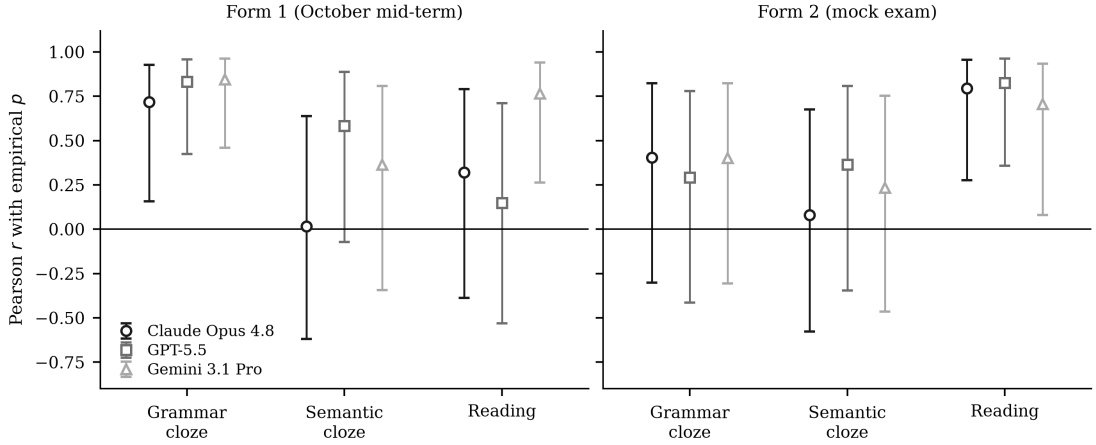


Figure 2: Validity by task type and form with 95% confidence intervals. On Form 1 every model is highest on grammar cloze; on Form 2 every model is highest on reading.

Form	Model	Δz	95% CI
Form 1	Claude Opus 4.8	0.89	[−0.16, 1.93]
	GPT-5.5	0.53	[−0.52, 1.57]
	Gemini 3.1 Pro	0.86	[−0.19, 1.90]
Form 2	Claude Opus 4.8	0.35	[−0.70, 1.40]
	GPT-5.5	−0.08	[−1.13, 0.97]
	Gemini 3.1 Pro	0.19	[−0.86, 1.24]

Table 3: Grammar versus semantic cloze. Δz = Fisher-transformed correlation for grammar minus semantic, within form and model. Descriptive: the three models share items and are not independent samples.

for reading, on a section whose empirical difficulties are themselves tightly clustered, so an easy item set can still carry between-item variation that the models track. The Form 1 reading coefficients, however, ranged from .15 to .77 across the three models, so these data do not support a general claim that reading difficulty is easier to predict than cloze difficulty.

Three of the four forms in this bank are compressed against the ceiling. They still yield precise item difficulties, and the reading result shows that a compressed form can carry usable signal, but they offer limited discrimination among higher-ability examinees, and a model whose predictions cluster in the moderately-easy range can show small absolute errors on them for reasons that have little to do with the quality of its judgements.

7 Limitations

Ten items per cell yield intervals wide enough that most contrasts in this paper are consistent with a wide range of true values. The three models predict the same items and are not independent replications. The data come from one cohort at a single school, so the results describe that testing program rather than L2 examinations in general. Form 2 item 50 was excluded because its options are pictures that were not supplied to the models.

8 Conclusion

On two forms of an operational Grade 9 English examination, three models produced difficulty es-

Form	Model	Bias	MAE	RMSE	SD ratio	Baseline
Form 1	Claude Opus 4.8	−.01	.145	**.175**	.53	.187
	GPT-5.5	+.15	.183	.226	.46	
	Gemini 3.1 Pro	+.17	.179	.229	.40	
Form 2	Claude Opus 4.8	−.16	.162	.198	1.91	.066
	GPT-5.5	−.03	.055	.073	1.01	
	Gemini 3.1 Pro	−.02	.047	**.064**	.89	

Table 4: Absolute calibration by form and model. Bias = mean predicted minus empirical *p*roportion correct. SD ratio = standard deviation of predictions divided by standard deviation of empirical difficulty. Baseline = RMSE of predicting the form mean for every item, a within-sample descriptive reference. Bold marks the only model that beats the baseline on each form.

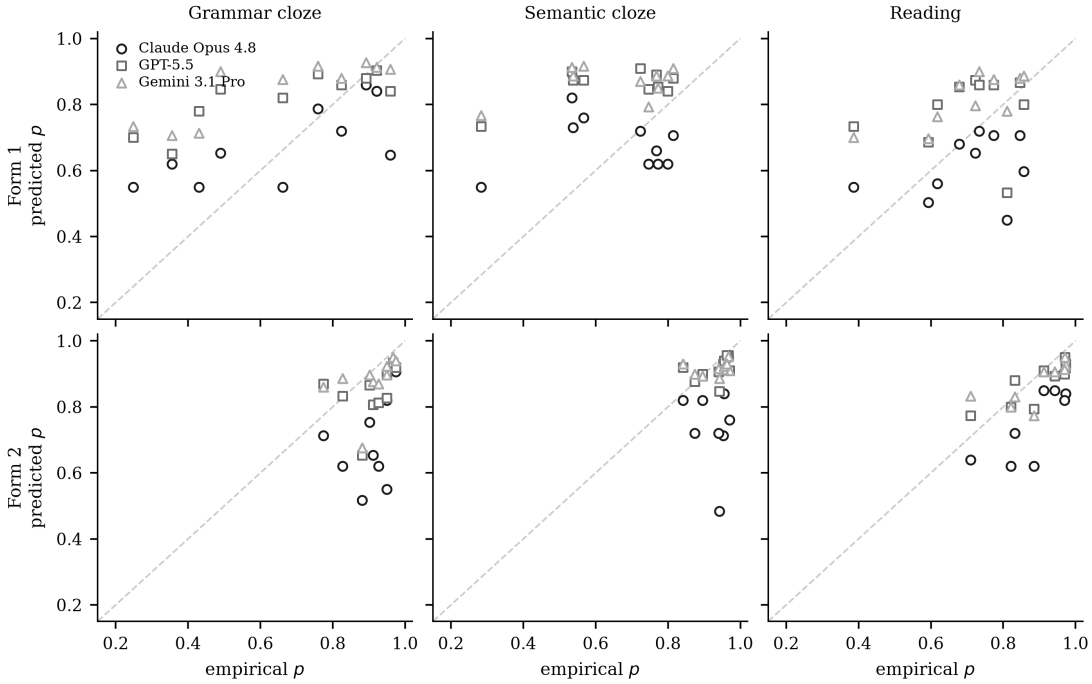


Figure 3: Predicted against empirical *p*roportion correct for every analysed item, by form and task type. The dashed line is perfect calibration. On Form 1 the predictions span a narrower range than the empirical values; on Form 2 the empirical values are themselves narrow and the relation between the two dispersions varies by model.

timates whose correlation with empirical difficulty varied by task type and by form: grammar cloze had the highest correlation for every model on the first form, reading for every model on the second. These data suggest that a task-type advantage observed on a small item set from one form needs checking on other forms before it is treated as a property of the model. The protocol, prompts and per-item prediction records are retained so that the estimates reported here can be checked and extended.

References

Luca Benedetto, Paolo Cremonesi, Andrew Caines, Paula Buttery, Andrea Cappelli, Andrea Giussani, and Roberto Turrin. 2023. A survey on recent approaches to question difficulty estimation from text. *ACM Computing Surveys*, 55(9):1–37.

Andrew Caines, Luca Benedetto, Shiva Taslimipoor, Christopher Davis, Yuan Gao, Øistein Andersen, Zheng Yuan, Mark Elliott, Russell Moore, Christopher Bryant, Marek Rei, Helen Yannakoudakis, Andrew Mullooly, Diane Nicholls, and Paula Buttery. 2023. On the application of large language models for language teaching and assessment technology. *Preprint*, arXiv:2307.08393.

Seonjeong Hwang, Hyounghun Kim, and Gary Geunbae Lee. 2025. Can LLMs estimate cognitive complexity of reading comprehension items? *Preprint*, arXiv:2510.25064.

Daria Kolesnikova, Kirill Fedyanin, Abe D. Hofman, Matthieu J. S. Brinkhuis, and Maria Bolsinova. 2026.

Estimating item difficulty with large language models as experts. *Preprint*, arXiv:2605.18562.

Vatsal Raina and Mark Gales. 2024. Question difficulty ranking for multiple-choice reading comprehension. *Preprint*, arXiv:2404.10704.

Pooya Razavi and Sonya Powers. 2025. Estimating item difficulty using large language models and tree-based machine learning algorithms. *Preprint*, arXiv:2504.08804.

Measuring Science Teachers' Generative AI Literacy for Classroom Assessment: Performance Assessment

Ruiping Huang and Yue Yin
College of Education, University of Illinois Chicago
1040 W Harrison St, Chicago, IL, 60607, USA

Abstract

We report preliminary analyses from an ongoing study using performance-based tasks to assess science teachers' generative AI (GenAI) literacy for classroom assessment and compare these scores with self-reported GenAI use and confidence. Preliminary findings indicate moderate correlation with use frequency but weak correlation with confidence, highlighting performance assessments' value.

Keywords: Generative AI; AI literacy; classroom assessment; performance assessment

1 Background

Widely available generative AI (GenAI) tools such as ChatGPT and Gemini have demonstrated substantial potential to strengthen classroom assessment, a critical component of instruction that provides evidence on students' strengths, needs, and learning progress to inform teachers' instructional decisions (Baidoo-Anu et al., 2025; Risko and Walker-Dalhouse, 2010). From item generation and automated scoring to personalized feedback and support for academic integrity (Naseer et al., 2024; Alier et al., 2024), GenAI is transforming how teachers design, administer, and interpret assessments.

However, teachers face substantial challenges in implementing GenAI, as evidenced by their limited knowledge, low confidence and interest, negative attitudes, ethical concerns, disruptions to traditional assessment practices, and difficulties in interpreting GenAI outputs (Yan et al., 2024; Alier et al., 2024; Nah et al., 2023; Celik et al., 2022; Zhai and Nehm, 2023). Effectively addressing these demands requires cultivating teachers' GenAI literacy (Yan et al., 2024; Nyaaba, 2024), particularly for classroom assessment.

It is critical to measure teachers' GenAI literacy specifically for classroom assessment using mixed-methods approaches for the following reasons. First, most existing AI literacy frameworks focus on students rather than teachers (Zhou et al., 2025; Lintner, 2024). Second, current assessments rely primarily on self-report surveys (Knoth et al., 2024; Ding et al., 2024). Very few studies used performance assessment, that is, how teachers use GenAI in teaching practice. They do not fully capture GenAI-specific competencies, such as prompt design, output evaluation, and ethical reasoning (Zhang and Magerko, 2025; Gu and Ericson, 2025). Third, most studies address general digital or AI literacy and are not tailored to the pedagogical and assessment demands of day-to-day classroom practice (Sparks et al., 2024).

2 Aims

To address these gaps, we developed performance assessment tasks to measure middle and high school science teachers' GenAI literacy for classroom assessment. Our study is designed to answer the following questions:

> (a) How reliable and valid are the performance assessment tasks?
>
> (b) How is the science teachers' GenAI literacy measured by the performance assessment tasks?
>
> (c) How is the science teachers' GenAI literacy measured by the performance assessment tasks compared with that measured by the self-reported survey?

3 Samples

We are currently collecting data. Twenty-one in-service middle and high school science teachers from U.S. Midwest school districts have been recruited and completed a self-reported survey. Table 1 presents the participants' characteristics and descriptive statistics.

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 448–458
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

Table 1: Participant Characteristics and Descriptive Statistics (N = 21)

Variable	Value	n (%)	M (SD)
Gender	Female	11 (52.38)	
	Male	10 (47.62)	
Race/Ethnicity	White	18 (85.71)	
	Hispanic/Latino	1 (4.76)	
	Asian	1 (4.76)	
	Other: Mexican American	1 (4.76)	
Education Level	Bachelor's degree	8 (38.10)	
	Master's degree	11 (52.38)	
	Doctoral degree	2 (9.52)	
Age Group	20–29	2 (9.52)	
	30–39	13 (61.90)	
	40–49	5 (23.81)	
	50-59	1 (4.76)	
School Level	Middle school	11 (52.38)	
	High school	10 (47.62)	
Subjects (multiple choices)	General Science	12 (57.14)	
	Biology	3 (14.29)	
	Chemistry	7 (33.33)	
	Physics	1 (4.76)	
	Earth/Space Science	3 (14.29)	
Student Access to Devices	Yes	21 (100)	
Teaching Experience (years)			11.29 (5.53)
Student IEP/504 (%)			21.00 (0.09)
English Learners (%)			20.00 (0.26)
Prior AI PD Participation hours			3.90 (8.55)
Self-Directed AI Learning Hours			5.57 (7.32)

Note. Categories not selected by any participant (e.g., Black, Pacific Islander) are omitted for clarity. The complete survey was in Appendix A.

Among them, nine teachers have completed the performance assessment tasks. We anticipate 30 science teachers will have completed both the self-reported survey and performance assessment tasks by the end of August 2026.

4 Methods

To guide the design of the instruments, we developed a conceptual framework (Figure 1) grounded in the iterative classroom assessment cycle (Bonner and Chen, 2019), which includes four phases: developing assessments, collecting responses, interpreting responses, and using results (Table 2). This cycle builds on foundational assessment literature (Buhagiar, 2006; Gottlieb and Katz, 2020; Ruiz-Primo and Furtak, 2007; Hajj-Hassan et al., 2014) and is refined to capture how teachers can engage with GenAI tools at each phase (Sipahioglu, 2024; Chaivisit et al., 2024; MacDowell et al., 2024).

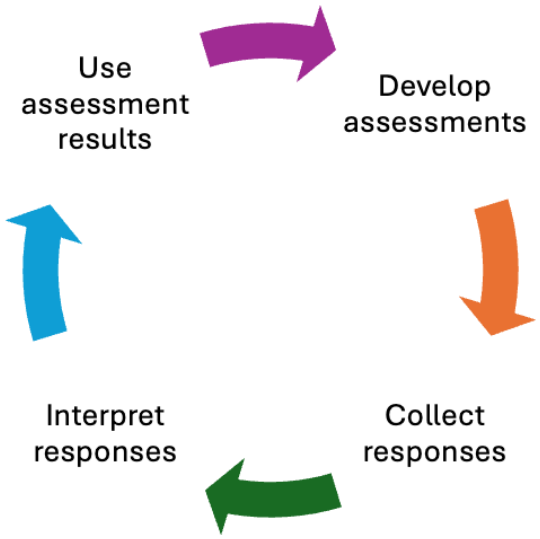


Figure 1: Teachers' classroom assessment cycle

At every phase of classroom assessment, GenAI can play a supportive role (Table 2). For example, in the assessment-development phase, GenAI can draft assessment items and tasks; in the response-collection phase, it can increase accessibility by adapting content and diversifying formats (e.g., visual or audio aids, translation); in the response-interpretation phase, it can assist with automated grading and organizing student work; and in the

use-of-results phase, it can suggest follow-up assessments and learning resources. We used this alignment between GenAI functions and assessment phases to develop instruments to measure teachers' GenAI literacy, specifically within the context of classroom assessment.

Table 2: GenAI-Supported Functions Across Phases of Teachers' Classroom Assessment

Phase	How GenAI Can Help Teachers for Classroom Assessment
Develop assessments	-Draft assessments (e.g. items, performance tasks) -Align assessments to standards/rubrics -Differentiate assessment materials by level or learner profile
Collect responses	-Make assessment inclusive and accessible by adapting content, diversifying formats, and reducing bias. -Simulate student responses to test clarity and difficulty -Transcribe/Translate student input when needed
Interpret responses	-Process and organize student responses (e.g., summarizing, tagging, identifying incomplete data) -Assist with automated grading
Use assessment results	-Draft feedback summaries and insights -Draft post-assessment reports, and action plans -Recommend follow-up assessments or learning resources

Appendix A is the performance assessment tasks we developed. Appendix B is the scoring rubric we developed to study teachers' GenAI literacy. Appendix C is the self-reported survey measuring GenAI literacy.

5 Anticipated Results

5.1 RQ1. Psychometric Features

We are still collecting data, refining the final performance assessment rubrics, and training coders. Table 3 shows the preliminary scoring results of the nine teachers. We found that the internal consistency of the four items was .78. Because the scoring rubrics and coder training are still being finalized, inter-rater reliability has not yet been established. We will report inter-rater reliability evidence in the full paper after rubric refinement and coder training are completed. The internal consistency of the confidence measure for the four phases are 0.89, 0.93, 0.95, and 0.95 respectively. The internal consistency of the frequency measure of the four phases are 0.86, 0.81, 0.93, and 0.79 respectively.

Table 3: Descriptive Statistics of Performance Assessment Scores (N = 9)

	Minimum	Maximum	Mean	Std. Deviation
Prompt	1.00	2.00	1.78	.44
Format and File Use	.00	2.00	1.00	.87
Communication	.00	2.00	.89	.93
Critical Evaluation	.00	2.00	1.00	.71
Total	.50	2.00	1.17	.57

5.2 RQ2. GenAI Literacy Profile

Preliminary analyses indicate the following patterns:

First, teachers vary in their proficiency with GenAI for classroom assessments on their performance assessment tasks (Table 3). As shown in Appendix D, teacher 6's prompt includes context, content, and output specifications, and instructional clarity, and scored 2, while teacher 1 includes only context and content specifications and scored 1 for the prompt.

Second, teachers' performance differs significantly across the four dimensions (Figure 2), $F(3, 24) = 4.59$, $p = 0.011$. Post hoc comparisons in a repeated ANOVA with the Sidak adjustment show that teachers, on average, scored significantly higher than the critical evaluation ($p = .004$).

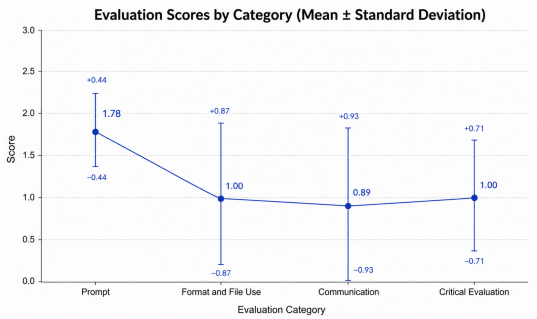


Figure 2: Scoring of the four dimensions

Third, we plan to further score and compare teachers' performance across dimensions to examine additional validity evidence. For example, teachers may score relatively high on one dimension, such as format and file use, but lower on another dimension, such as critical evaluation or communication. Such patterns can help identify the aspects of GenAI literacy that may require additional instructional or professional development support.

5.3 RQ3. Relationship with Self-reported Scales

Preliminary results show that teachers' performance scores are moderately correlated with their

frequency of GenAI use, r = .63, p = .07; but the correlation between their self-rated confidence and performance score is relatively weak, r = .25, p =.52.

We will further compare teachers' scores on the performance assessment tasks across four assessment phases and their self-reported frequency of GenAI use and GenAI confidence.

6 Anticipated Conclusions

Our study will fill the following gaps: First, instruments are needed to measure teachers' AI literacy for classroom assessment. Second, performance assessments are lacking to measure AI literacy, while most scales are self-reported surveys.

Our performance assessment tasks will help better understand teachers' current use of GenAI in their assessment practice. Although the current performance assessment tasks were designed for middle and high school science teachers, the classroom assessment framework and scoring rubrics are general rather than subject-specific. Thus, the tasks can be adapted for teachers in other subject areas or educational contexts by revising the classroom scenarios, disciplinary content, and expected assessment products.

The comparison between teachers' performance assessment and their self-reported data will help inform the validity study of measurements.

These findings should be interpreted as preliminary. Data collection is ongoing, and the full paper will provide an updated sample description, report inter-rater reliability after rubric refinement and coder training, and examine additional validity evidence through relationships among the performance dimensions as the available data allow.

References

Marc Alier, Francisco José García Peñalvo, and Jorge D. Camba. 2024. Generative Artificial Intelligence in Education: From Deceptive to Disruptive. *International Journal of Interactive Multimedia and Artificial Intelligence*, 8(5):5–14.

David Baidoo-Anu, Amirhossein Rasooli, Christopher DeLuca, and Liying Cheng. 2025. Conceptions of classroom assessment and approaches to grading: teachers' and students' perspectives. *Education Inquiry*, 16(3):411–439.

Sarah M. Bonner and Peggy P. Chen. 2019. *Systematic Classroom Assessment: An Approach for Learning and Self-Regulation*, 1 edition. Routledge.

Michael A. Buhagiar. 2006. The classroom assessment cycle within the alternative assessment paradigm: exploring the role of the teacher.

Ismail Celik, Muhterem Dindar, Hanni Muukkonen, and Sanna Järvelä. 2022. The Promises and Challenges of Artificial Intelligence for Teachers: a Systematic Review of Research. *TechTrends*, 66(4):616–630.

Sarinporn Chaivisit, Tutaleni Iita Asino, Sukanda Jongsermtrakoon, Penny Thompson, Fatemeh Rezaie, and Sutithep Siripipattanakul. 2024. Empowering Educators With Generative AI Tools and Support:. In *Advances in Educational Technologies and Instructional Design*, pages 56–81. IGI Global.

Lu Ding, Sohee Kim, and R. Allan Allday. 2024. Development of an AI literacy assessment for non-technical individuals: What do teachers know? *Contemporary Educational Technology*, 16(3):ep512.

Margo Gottlieb and Anne Katz. 2020. Assessment in the Classroom. In *The Encyclopedia of Applied Linguistics*, pages 1–8. Wiley.

Xingjian Gu and Barbara J. Ericson. 2025. AI Literacy in K-12 and Higher Education in the Wake of Generative AI: An Integrative Review. *arXiv preprint*.

Mohamad Hajj-Hassan, Bassam Hussein, Khaled Chahine, and Amin Haj-Ali. 2014. A Sustainable Approach for Program Students Outcomes Assessment: The Case of the Department of Biomedical Engineering at the Lebanese International University. *Journal of Modern Education Review*, 4(7):555–563.

Nils Knoth, Marie Decker, Matthias Carl Laupichler, Marc Pinski, Nils Buchholtz, Katharina Bata, and Ben Schultz. 2024. Developing a holistic AI literacy assessment matrix – Bridging generic, domain-specific, and ethical competencies. *Computers and Education Open*, 6:100177.

Tomáš Lintner. 2024. A systematic review of AI literacy scales. *npj Science of Learning*, 9(1):50.

Paula MacDowell, Kristin Moskalyk, Katrina Korchinski, and Dirk Morrison. 2024. Preparing Educators to Teach and Create With Generative Artificial Intelligence. *Canadian Journal of Learning and Technology*, 50(4):1–23.

Fiona Fui-Hoon Nah, Ruilin Zheng, Jingyuan Cai, Keng Siau, and Langtao Chen. 2023. Generative AI and ChatGPT: Applications, challenges, and AI-human collaboration. *Journal of Information Technology Case and Application Research*, 25(3):277–304.

Fawad Naseer, Muhammad Usama Khalid, Nafees Ayub, Akhtar Rasool, Tehseen Abbas, and Muhammad Waleed Afzal. 2024. Automated assessment and feedback in higher education using generative AI. In *Transforming education with generative AI:*

Prompt engineering and synthetic content creation, pages 433–461. IGI Global Scientific Publishing.

Matthew Nyaaba. 2024. Transforming Teacher Education in Developing Countries: The Role of Generative AI in Bridging Theory and Practice. *arXiv preprint*.

Victoria J. Risko and Doris Walker-Dalhouse. 2010. Making the Most of Assessments to Inform Instruction. *The Reading Teacher*, 63(5):420–422.

Maria Araceli Ruiz-Primo and Erin Marie Furtak. 2007. Exploring teachers' informal formative assessment practices and students' understanding in the context of scientific inquiry. *Journal of Research in Science Teaching*, 44(1):57–84.

Mete Sipahioglu. 2024. Empowering Teachers With Generative AI Tools and Support:. In *Advances in Educational Technologies and Instructional Design*, pages 214–238. IGI Global.

Jesse R. Sparks, Teresa M. Ober, Caitlin Tenison, Burcu Arslan, Ido Roll, Paul Deane, Diego Zapata Rivera, Reginald M. Gooch, and Tenaha O'Reilly. 2024. Opportunities and challenges for assessing digital and ai literacies. Technical report, ETS Research Institute.

Lixiang Yan, Samuel Greiff, Ziwen Teuber, and Dragan Gašević. 2024. Promises and challenges of generative artificial intelligence for human learning. *Nature Human Behaviour*, 8(10):1839–1850.

Xiaoming Zhai and Ross H. Nehm. 2023. AI and formative assessment: The train has left the station. *Journal of Research in Science Teaching*, 60(6):1390–1398.

Chengzhi Zhang and Brian Magerko. 2025. Generative AI Literacy: A Comprehensive Framework for Literacy and Responsible Use. *arXiv preprint*.

Xinyan Zhou, Yan Li, Ching Sing Chai, and Thomas K. F. Chiu. 2025. Defining, enhancing, and assessing artificial intelligence literacy and competency in K-12 education from a systematic review. *Interactive Learning Environments*, 33(10):5766–5788.

Appendix A. Performance Assessment Tasks

Please use the ChatGPT account we created for you to complete the following tasks.

Website: https://chatgpt.com/
Account: xxx
Password: xxx

Task 1: Develop Assessments

- You are preparing an assessment at the end of a science unit. Please use ChatGPT to generate a classroom assessment to measure students' learning outcome and its corresponding solutions.

Task 2: Collect responses

- Please use ChatGPT to create sample students' responses to an existing science assessment from your science class.

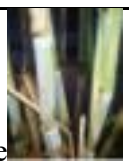
Task 3: Interpret Responses

- Suppose the following students are your real students in your classroom, below are their responses to a science question (comparing types of natural sugars). Use ChatGPT to help you score/analyze the students' responses.

Assessment Tasks for Students: Comparing Types of Natural Sugars

Amy wants to know if the sugar found in honey, milk, sugarcane, and apples are the same. To find this out, she chemically removed the sugar from each food and recorded the properties of the sugar in Table 1.

Table 1: Characteristic properties of sugar in different foods

Source of sugar sample	Density	Solubility in Water	Melting Point
Honey 	1.69 g/cm^3	Yes	103 °C
Milk 	1.53 g/cm^3	Yes	202 °C
Sugarcane 	1.59 g/cm^3	Yes	186 °C
Apple 	1.70 g/cm^3	Yes	103 °C

Which of the above sources has the same type of sugar? Why?

Right Answer:

Apple and honey have the same type of sugar, as they share the same density (with any small differences likely due to measurement error) and melting point. Both density and melting point are characteristic properties used to identify substances.

Here are 6 students' responses:

Alex Martinez: Well the one that has the similar sugar is honey and apple because there only 1 percent from each other. Honey and apple, because there like 1 sugar away from each other. And milk and sugar are 6 percent of sugar away from each other.
Emily Chen: My conclusion is that none have the same sugar and two of them have the same melting point and they all have the same solubility in water. Differences are density also sugar can and milk are different in melting point from honey and apple. Honey and apple have the same melting point. All of them (honey,milk,sugar cane, apple) have the same solubility in water. They all have a different density. Milk's and sugar cane's melting point are different from honey's and apple's melting point.
Jordan Williams: A lot of foods have the same sugar but some foods have a decent amount of sugar the differences is that sugar can be in anything the simularites is that they have foods that are the same with sugar The similarities of the foods is that some of the food have more sugar than the other kinds of food and some people put the same sugar in foods and put water in them for the differences
Sophia Nguyen: What i can conclude from the pattern is that the apple and honey sugar are almost the most because they have the same melting point and they almost have the amount of density. The similarities or differences i would look to tell weather they have the same sugar is the Density and the melting point.
Michael Johnson: I can draw a conclusion that the sugar of the sources from both the honey and the apple are the same because their characteristic properties are either the same or very, very similar. For example, the density of the honey is 1.69 g/cm cubed, and the apple's density is 1.70g/cm cubed, so the difference between both the densities is only 0.01g/cm cubed. Additionally, their are both soluble and have the melting points, which is 103 degrees Celsius. However, the only thing similar between the milk and the sugar cane is that they are both soluble, and other than that, the density of each are not the same, and they have two, very different melting points, meaning that they are most likely not made of the same sugar. In order to tell whether or not any of the foods contained the same sugar, I would look to see if the density and melting point were the same or, at least, similar. Also, I would make sure that the solubility of the sugar from the different sources are the same. In conclusion, I would make sure that the characteristic properties between the sugar from the multiple sources were the same because the properties are what make up a substance and what makes each one similar or different, so if all of the properties for each source are the same, then it would mean that the sugar used in each sample was also the same.
Kayla Thompson: Claim: honey and apple have the same sugar. evidence: the density (1.69 and 1/70 g/cm3) and melting point are the same (103 C). Reasoning: unique substances that have unique characteristic properties.

Task 4: Use Assessment Results

- Building on the analysis results in Task 3, please use ChatGPT to support teaching and learning based on these assessment results.

Appendix B. Performance Assessment Scoring Rubrics

Criterion	Definition (What is being evaluated)	0 = Novice	1 = Emerging	2 = Experienced
Prompt Construction	How clearly and effectively the teacher constructs prompts to guide GenAI in achieving the intended task goals. Prompts are evaluated based on the presence of four core prompt features: (a) Context (e.g. grade level, student characteristics); (b) Content (such as plate tectonics as topic or content) (c) Output (Question format, structure, length, style, or inclusion of examples); (d) Instructional clarity (clear, complete, and actionable instruction, explicit guidance such as step-by-step directions, role assignment, or explicit constraints or criteria, such as align with Indiana standard)	Demonstrates none of the core prompt features. Prompts largely copy or restate the original task instructions with little or no modification.	Demonstrates 1–2 of the core prompt features across the four tasks (features may appear separately across prompts or in combination within a single prompt).	Demonstrates 3-4 core prompt features across the four tasks (features may appear separately across prompts or in combination within a single prompt). Prompts are aligned with task goals.
Format and File Use	The extent to which the teacher can appropriately use different input and output formats with GenAI to support task completion, including requesting responses in useful forms and working with uploaded materials.	Uses GenAI primarily in plain text form. Does not request specific output formats (e.g., tables, worksheets, doc, pdf, images, canvas-based drafts) and does not use file upload features.	Uses 1–2 distinct formats or file interactions (e.g., requests a table or worksheet, or uploads a PDF/document/image) or attempts format use 1–2 times.	Purposefully selects and uses appropriate formats to support the task (e.g., asks GenAI to organize content into tables, worksheets, structured documents, or canvas-based drafts) and effectively works with uploaded materials when relevant (e.g., PDFs, images). Demonstrates ≥3 instances of format or file use.
Interactive Communication (B, C)	The extent to which the teacher engages in ongoing, responsive communication with GenAI during task completion, using follow-up interactions to clarify, refine, or redirect responses.	Accepts initial GenAI outputs with no follow-up interaction or clarification.	Engages in 1-2 follow-up interactions (e.g., minor revisions or clarifications), but these interactions do not improve subsequent GenAI responses.	Actively communicates with GenAI through ≥3 purposeful follow-up requests, clarifications, or redirections informed by prior outputs.
Critical Evaluation and Decision-Making (C, E) (Sullivan, Michael, et al., 2024)	The extent to which the teacher critically evaluates GenAI outputs (e.g., accuracy, appropriateness) and uses professional judgment to accept, revise, or reject outputs to make task-relevant decisions.	Accepts outputs without verification or critique.	Identifies issues in the GenAI output (e.g., inappropriate content, unclear format) but conducts limited evaluation or doesn't follow up or give up easily.	Not only identifies issues, but also actively verifies or challenges outputs (e.g., checks sources, requests evidence, or asks GenAI to justify or revise responses).

Note: B = Behavioral; C = Cognitive; E = Ethical

Appendix C. Self-reported Scales to Measure Teacher's GenAI Literacy

Phase 1: Develop Assessments

- How frequently did you use AI for each of the following assessment activities in spring 2025? (Left)
- How confident are you now in using AI for each of the following assessment activities? (Right)

	My Frequency in Spring 2025					My Current Confidence Level				
	Never (1)	Rarely (once a semester) (2)	Sometimes (several times a semester) (3)	Often (monthly) (4)	Always (weekly) (5)	Strongly unconfident (1)	Unconfident (2)	Neutral (3)	Confident (4)	Strongly confident (5)
Use AI to modify existing classroom assessments.	○	○	○	○	○	○	○	○	○	○
Use AI to draft classroom assessments (e.g. items and performance-based tasks).	○	○	○	○	○	○	○	○	○	○
Use AI to develop scoring rubrics or standards for classroom assessments.	○	○	○	○	○	○	○	○	○	○
Use AI to develop solutions to answers for classroom assessments.	○	○	○	○	○	○	○	○	○	○
Use AI to design classroom assessments that require students to engage in their own thinking instead of relying solely on AI.	○	○	○	○	○	○	○	○	○	○

Phase 2: Collect Responses

- How frequently did you use AI for each of the following assessment activities in spring 2025? (Left)
- How confident are you now in using AI for each of the following assessment activities? (Right)

	My Frequency in Spring 2025					My Current Confidence Level				
	Never (1)	Rarely (once a semester) (2)	Sometimes (several times a semester) (3)	Often (monthly) (4)	Always (weekly) (5)	Strongly unconfident (1)	Unconfident (2)	Neutral (3)	Confident (4)	Strongly confident (5)
Use AI to translate and administer classroom assessment to accommodate English learners.	○	○	○	○	○	○	○	○	○	○
Use AI to modify and administer classroom assessments in accessible formats for learners with special needs (e.g., audio and visual versions).	○	○	○	○	○	○	○	○	○	○
Use AI to revise and administer classroom assessment materials to accommodate students with different culture backgrounds.	○	○	○	○	○	○	○	○	○	○
Use AI to simulate student responses to test the clarity and difficulty of classroom assessments.	○	○	○	○	○	○	○	○	○	○
Use AI to transcribe	○	○	○	○	○	○	○	○	○	○

student responses when needed.										
Use AI to translate student responses when needed.	○	○	○	○	○	○	○	○	○	○

Q20 Phase 3: Interpret Responses

- How frequently did you use AI for each of the following assessment activities in spring 2025? (Left)
- How confident are you now in using AI for each of the following assessment activities? (Right)

	My Frequency in Spring 2025					My Current Confidence Level				
	Never (1)	Rarely (once a semester) (2)	Sometimes (several times a semester) (3)	Often (monthly) (4)	Always (weekly) (5)	Strongly unconfident (1)	Unconfident (2)	Neutral (3)	Confident (4)	Strongly confident (5)
Use AI to summarize student responses to classroom assessments.	○	○	○	○	○	○	○	○	○	○
Use AI to score students' responses to objective classroom assessments (e.g., multiple-choice, true/false, matching).	○	○	○	○	○	○	○	○	○	○
Use AI to score students' responses to open-ended classroom assessments.	○	○	○	○	○	○	○	○	○	○
Use AI to analyze students' responses to classroom assessments.	○	○	○	○	○	○	○	○	○	○
Use AI to group students based on their responses to classroom assessments.	○	○	○	○	○	○	○	○	○	○

Phase 4: Use Assessment Results

- How frequently did you use AI for each of the following assessment activities in spring 2025? (Left)
- How confident are you now in using AI for each of the following assessment activities? (Right)

	My Frequency in Spring 2025					My Current Confidence Level				
	Never (1)	Rarely (once a semester) (2)	Sometimes (several times a semester) (3)	Often (monthly) (4)	Always (weekly) (5)	Strongly unconfident (1)	Unconfident (2)	Neutral (3)	Confident (4)	Strongly confident (5)
Use AI to draft feedback for students based on classroom assessment results.	○	○	○	○	○	○	○	○	○	○
Use AI to recommend follow-up assessments or learning resources based on classroom assessment results.	○	○	○	○	○	○	○	○	○	○
Use AI to reflect on student learning progress using classroom assessment results.	○	○	○	○	○	○	○	○	○	○
Use AI to help make instructional decisions based on classroom assessment results to assignments.	○	○	○	○	○	○	○	○	○	○
Use AI to support my communication with parents on classroom assessment results.	○	○	○	○	○	○	○	○	○	○

Appendix D. Examples of Teachers' Prompts in Phase 1. Assessment Generation

- Teacher 6: "*You are an experienced chemistry teacher with a masters degree in chemistry. You are teaching an honors chemistry class for advanced sophomores that will get dual credit at completion of the course through a community college* **[context specification]**. *I would like you to make an end of unit assessment to measure student learning on kinetics and equilibrium* **[content specification]** *that is between 20 and 25 questions that must include questions about an ice table, questions on le chatleier principle, and a problem where students must determine rate order from lab data* **[output specification]**. *I would like you to also create a key for the assessment. The assessment must fit within indiana state standards in chemistry and A typical Chemistry 101 and 102 course at the university level* **[Instructional clarity]**."

- Teacher 1: "*Generate an end-of-unit assessment testing high school students'* **[context specification]** *knowledge of plate tectonics, specifically the three main types of plate boundaries as well as intraplate activities* **[content specification]**."

Do Synthetic Students Thrive? Validity and Fairness of Generated Assessment Data

Neba Nfonsang and Temple Lovelace
Assessment for Good - AERDF and Oluko Learning
{nnfonsang, tlovelace}@aerdf.org

Abstract

Synthetic assessment data can match real distributions while losing the within-group spread or correlation structure that asset-based instruments reveal in learners furthest from opportunity. Using a stratified, majority Black and Latinx sample of students, the authors evaluated whether synthetic assessment data generated through a retrieval-augmented generation (RAG) pipeline preserved the construct structure of ground-truth assessment data. Synthetic responses were evaluated using both distributional and psychometric measures, including response distributions, response agreement, confirmatory factor analysis, and model fit indices. The results of this study demonstrate that synthetic assessment data can closely reproduce the distributions of the original data. However, limitations were found. These data did not preserve underlying correlation structures and psychometric properties, highlighting the limitations of relying solely on distributional similarity when evaluating synthetic assessment data. These findings suggest that psychometric preservation should be considered alongside distributional similarity when evaluating synthetic assessment data.

1 Introduction

1.1 Background

Synthetic data is increasingly used to validate measures, enlarge small samples, and share data without exposing real study participants. According to Kalkbrenner (2021), social scientists contribute to research by developing and validating scores on instruments as well as evaluating the psychometric properties of existing measures for use with research participants.

Synthetic data are defined as artificially generated data that are created to resemble the characteristics of real data without containing actual records from real individuals. Educational measurement research relies on access to high-quality assessment data to develop, validate, and refine instruments. While simulated datasets are not new in assessment research, particularly in educational measurement, recent advances in artificial intelligence and machine learning have renewed interest in the use of synthetic data. These approaches enable researchers to generate realistic datasets for instrument development, validation, methodological research, and privacy-preserving data sharing.

Researchers often face limitations related to sample size, privacy restrictions, incomplete responses, and the underrepresentation of certain populations, especially historically underrepresented populations. Inadequate sample sizes can hinder efforts to establish the psychometric quality and validity of assessment instruments (Anthoine et al., 2014).

Most recently, in the wake of frontier model use and generative artificial intelligence, synthetic data have continued to grow as a promising approach for addressing these challenges because synthetic data can augment limited datasets, facilitate data sharing while preserving privacy as well as support

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 459–463
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

experimentation with new assessment items, predictive modeling, methodological and machine learning research (İpek et al., 2026).

According to Jiang et al. (2024), synthetic data has potential to enable proprietary data access for statistical and research purposes and increase available data. The authors highlight the promise of synthetic data in enabling data sharing across institutions while preserving patient privacy and illustrate its utility for data augmentation and research.

Even more, Liu et al. (2025), stated that, educational researchers and practitioners can leverage synthetic data to release data without compromising student privacy, thereby promoting the development of open science and contributing to the advancement of education research. They also emphasized that high-volume, high-quality and diverse datasets are crucial for advancing research in the education field.

Finally, Vie, Rigaux & Minn (2022) argued that synthetic data generation opens new opportunities for research in education. In their paper, they present a generative model for educational data and an evaluation framework for comparing synthetic data generators. According to these authors, the comparison between real and synthetic datasets involved examining response distributions through histograms and evaluating measurement properties using Item Response Theory (IRT) estimates derived from both the real and generated data.

1.2 Problem Statement

Despite the advantages of synthetic data, it remains unclear whether synthetic data preserve the construct structure and learner variability captured by assessment instruments. Most evaluations of synthetic data focus on distributional fidelity, assessing the extent to which generated data reproduce the statistical characteristics of the original dataset. However, the psychometric properties of synthetic data may not match those of the original data for different reasons. For example, data generator functions are not neutral instruments.

Additionally, synthetic data generators, including retrieval-augmented generation (RAG) systems and large language models, are influenced by their underlying model architectures, training data, knowledge bases, and retrieval mechanisms. As a result, the synthetic data they produce partly reflect the assumptions embedded within the generation process. This raises important concerns for asset-based assessment, particularly among learners from historically marginalized communities.

While asset-based instruments are designed to capture the diversity of learners' strengths and durable skills, synthetic generation processes may unintentionally reduce or homogenize variation within groups. For asset-based measurement of learners furthest from opportunity, the thing most at risk is within-group heterogeneity. Deficit-oriented systems compress these learners toward a single below-standard score, while the instrument and co-created construct ontology behind this work were built with learners to show the opposite, that they sit across the full range of durable-skill assets. We ask whether a synthetic generator keeps that range or flattens it, a question the field must contend with and one with direct consequences for equity.

Hence, an important question is whether synthetic data preserve the within-group heterogeneity observed among real learners. The answer has important implications for equity and fairness in assessment. Equally important is whether synthetic data preserve the construct structure reflected in learner response patterns.

1.3 Aims

This study makes three contributions. First, we present a retrieval-augmented generation (RAG) synthetic data pipeline developed for durable-skills assessment and report a finding about the role of its language-model component. Second, we propose diagnostics for evaluating whether a synthetic data generator preserves construct structure and within-group heterogeneity rather than pulling learners toward group-level averages. Third, we argue that standard algorithmic fairness approaches, which often seek to remove demographic information from models, may not be the appropriate framework for asset-based measurement.

2 Samples and Methods

2.1 Samples

The analyses used a durable-skills instrument that the authors created, administered to a sample that is majority Black and Latinx by deliberate design. The dataset consisted of 4,279 unique students who provided responses to a durable-skills assessment comprising 100 items administered to lower-grade

students and 130 items administered to higher-grade students. Across all students and items, the dataset contained approximately 339,000 item responses, yielding a large response matrix suitable for the RAG knowledge base. The ground-truth dataset used for the synthetic data evaluation consisted of 3,225 unique students who completed both the existing assessment items included in the RAG knowledge base and the newly developed items scored by the RAG system. Items were scored on 3-point, 4-point, or 5-point Likert-type scales. Assessment administration and validation activities were conducted under Institutional Review Board (IRB) approval.

2.2 Methods

Three generation approaches were used: multiclass classifiers (logistic regression, random forest, XGBoost) for generating scores of students on existing items, a retrieval or RAG system for new items, and a variational autoencoder for synthetic profiles. The retrieval finding is the methodological center of this paper. Figure 1 illustrates the architecture of the RAG pipeline. RAG is a technique used in modern AI systems to enrich prompts sent to large language models (LLMs). RAG enables LLMs to retrieve and incorporate information from specific internal knowledge bases, allowing LLMs to generate responses that are contextually relevant.

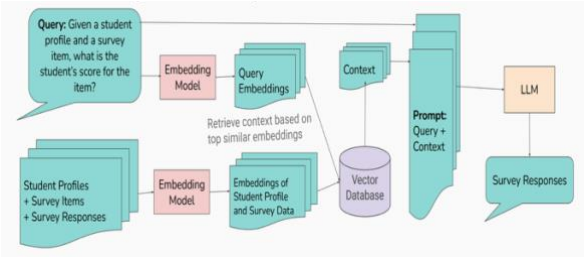


Figure 1: Retrieval-Augmented Generation (RAG) pipeline.

The RAG workflow was developed and executed within a secure, access-controlled private sandbox environment. Embedding generation, retrieval, and language model inference were performed entirely within the isolated environment, ensuring that student-level data, embeddings, retrieved context, and prompts remained local and were not transmitted to any external model provider or third-party service.

The RAG workflow involves converting student profiles, items, and responses into documents, storing responses as metadata to avoid conditioning the context on response values, generating embeddings from documents, and storing them in a vector database. A retriever then identifies relevant context, which is combined with the query in a prompt and sent to the LLM to generate a response. We implemented hierarchical indexing to improve retrieval performance by clustering documents and retrieving context in multiple stages.

When we asked the language model to explain its response selections during retrieval, it was almost always choosing the most common response among the retrieved students, behavior consistent with a nearest-neighbor approach. We therefore replaced the LLM component of the RAG pipeline with an explicit K-nearest-neighbor (KNN) algorithm, implemented as both a classifier and regressor, which reproduced the LLM's outputs in 97.2% of cases. Replacing the LLM with a KNN classifier and regressor reduced generation time from approximately 1,000 responses per hour to 1,000 responses in under five minutes, while maintaining 97.2% agreement with the RAG-LLM outputs.

We assessed the quality of the synthetic data using distributional and psychometric metrics to determine whether the generated data preserved within-group variability and construct structure.

3 Results

Since the RAG pipeline was used to generate student responses, we evaluated the quality of the synthetic data generated by the RAG pipeline by comparing synthetic and ground-truth response distributions, assessing the agreement between generated and observed responses, and examining whether the synthetic data preserved the psychometric properties and measurement characteristics of the ground truth dataset.

3.1 Evaluation of Distributions

The RAG-generated responses exhibited a distribution that closely matched the distribution of the actual responses as shown in Figure 2 (Grades 3-5).

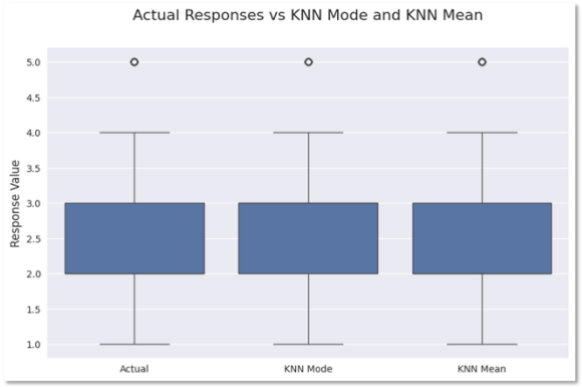


Figure 2: Distribution of actual and RAG-generated responses.

3.2 Evaluation of Accuracy

Agreement between generated and observed responses was assessed using exact match accuracy and accuracy with predictions falling within ±1 response category as shown in Table 1.

Grade Band	Method	Exact Match	Accuracy Within ±1
Lower Grades	KNN Mode	37.1%	83.3%
Lower Grades	KNN Mean	38.4%	90.6%
Higher Grades	KNN Mode	37.0%	82.6%
Higher Grades	KNN Mean	38.7%	90.5%

Table 1: Exact match accuracy and accuracy within ±1.

KNN Mean consistently outperformed KNN Mode across both grade bands, achieving approximately 38% exact-match accuracy and 90% accuracy within one response category.

3.3 Evaluation of Psychometric Properties

Although the RAG approach performed reasonably well based on accuracy within ±1, the generated responses often failed to preserve the correlation structure among items. The confirmatory factor analysis (CFA) revealed substantial differences between the RAG-generated and ground-truth data. As shown in Figure 3 (Grades 3-5), the RAG-generated constructs exhibited weak and, in some cases, negative factor loadings. Consistent with these results, RMSEA values indicated good fit for the ground-truth constructs (0.034-0.047) but poor fit for the RAG-generated constructs (0.138-0.173), suggesting that the synthetic data failed to preserve the original construct structure.

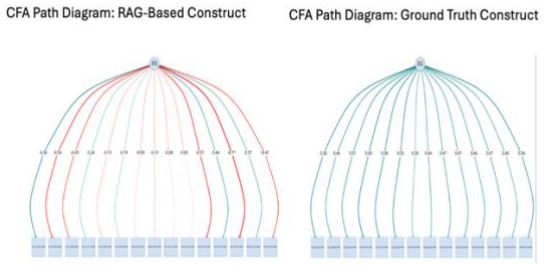


Figure 3: Standardized factor loading of RAG-based and ground truth constructs.

3.4 Variational Autoencoder Synthetic Data Evaluation

Variational Autoencoders (VAEs) generate synthetic data by reconstructing the underlying patterns in the original data. The psychometric properties of the VAE-generated assessment responses more closely resembled those of the ground-truth data than did the RAG-generated responses, as shown in Figure 4. Consistent with these results, RMSEA values indicated good fit for the ground-truth constructs (0.034-0.047), whereas the six VAE-generated constructs yielded RMSEA values ranging from 0.087 to 0.117. Although the VAE-generated data did not fully reproduce the fit observed in the ground-truth models, they demonstrated substantially better fit than the RAG-generated constructs (0.138-0.173), indicating stronger preservation of the original construct structure.

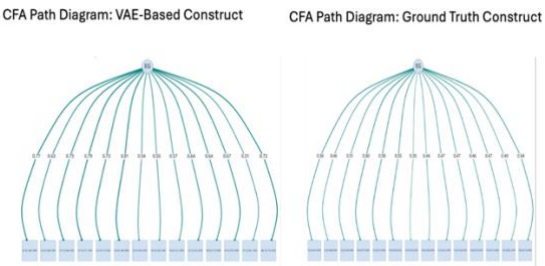


Figure 4: Standardized factor loading of VAE-based and ground truth constructs.

The distributions of the VAE-generated student profiles closely resembled those of the ground truth data. The VAE synthetic data demonstrated high similarity to the original dataset, achieving a 92.9% column shapes score, indicating strong preservation of individual variable distributions, and an 86.2% column pair trends score, indicating

that relationships between variables were largely maintained. These metrics yielded an overall quality score of 89.6%, calculated as the average of the column shapes and column pair trends scores, indicating strong preservation of both variable-level distributions and inter-variable relationships. Unlike the RAG-generated responses, the VAE-generated data preserved both the statistical distributions and psychometric structure of the original dataset, making them suitable for privacy-preserving data sharing and data augmentation applications.

4 Conclusions

We offer three contributions. First, we present a working synthetic-assessment pipeline and a finding about when a language model is doing no more than nearest-neighbor selection.

Second, we treat the generator as a context-laden tool rather than a neutral instrument, and we argue that the standard fairness move of removing demographic information may not be an appropriate frame for asset-based measurement, because it treats a demographic-neutral default as fair and erases the context the instrument exists to honor.

Third, we introduce a set of evaluation diagnostics that extend beyond distributional similarity. Our results suggest that preserving statistical distributions alone is insufficient. Although the RAG synthetic data closely resembled the ground-truth data in their overall distributions, they failed to preserve important correlation structures and psychometric properties.

These findings highlight the need to evaluate synthetic data not only on distributional similarity but also on their ability to preserve meaningful measurement characteristics, construct structure, and learner variability. Synthetic assessment data should therefore be evaluated not only on their ability to reproduce observed distributions, but also on their ability to preserve measurement structures that support valid inferences about learners.

This study is limited to six durable-skills instruments. Although the instruments span multiple durable-skills constructs, the RAG-related findings reported here may not generalize to other assessment domains, populations, foundation models, or synthetic data generation approaches. Moreover, the present results do not establish that the generated data are safe for public release, unrestricted sharing, or downstream data augmentation.

Future research should examine whether the tendency of RAG-generated assessment data to preserve response distributions while failing to preserve construct structure is observed across other educational assessments and synthetic data generation approaches.

References

Anthoine, E., Moret, L., Regnault, A., Sébille, V., & Hardouin, J.-B. (2014). Sample size used to validate a scale: A review of publications on newly-developed patient reported outcomes measures. *Health and Quality of Life Outcomes, 12*, 176. https://doi.org/10.1186/s12955-014-0176-2

İpek, Z. H., Poyraz, E., & Yiğit, G. (2026). The role of synthetic data in educational research: A systematic review. *International Journal of Interactive Mobile Technologies, 20*(16). https://doi.org/10.3991/ijim.v20i16.62224

Jiang, Y., García-Durán, A., Losada, I. B., Girard, P., & Terranova, N. (2024). Generative models for synthetic data generation: Application to pharmacokinetic/pharmacodynamic data. *Journal of Pharmacokinetics and Pharmacodynamics, 51*(6), 877–885. https://doi.org/10.1007/s10928-024-09935-6

Kalkbrenner, M. T. (2021). A practical guide to instrument development and score validation in the social sciences: The MEASURE approach. *Practical Assessment, Research, and Evaluation, 26*(1), Article 1. https://doi.org/10.7275/svg4-e671

Liu, Q., Shakya, R., Jovanovic, J., Khalil, M. I., & De La Hoz Ruiz, J. (2025). Ensuring privacy through synthetic data generation in education. *British Journal of Educational Technology, 56*(3), 1053–1073. https://doi.org/10.1111/bjet.13576

Vie, J.-J., Rigaux, T., & Minn, S. (2022). Privacy-preserving synthetic educational data generation. *arXiv*. https://arxiv.org/abs/2207.03202

AI-Enabled Quality Assurance for Multiple-Choice Assessment Items

Steven Moore
George Mason University
stevenjamesmoore@gmail.com

Nicholas Diana
Colgate University
ndiana@colgate.edu

Abstract

Generating multiple-choice questions is increasingly scalable, but establishing their assessment quality remains difficult. We present a focused narrative review of automated item-writing flaw detection, revision, psychometric screening, and NLP benchmark auditing. Database searches, citation retrieval, and nominated sources yield fourteen research reports reviewed in full text. We distinguish surface checks from content-sensitive judgments and map a 19-criterion rubric to detection methods and reported evidence. High label-level accuracy often coexists with weak positive-case detection, while rubric definitions and reference standards vary. Revision evidence is mixed, and the associations reported in prior work do not establish the effects of repair. We propose evaluating quality assurance as a sequence of independently validated decisions, with criterion-specific reporting, calibrated human review, and outcome-based assessment of revisions.

1 Introduction

Multiple-choice questions (MCQs) support efficient assessment but require careful design (Gottlieb et al., 2023). A distractor may be grammatically inconsistent, implausible, or defensible as a second answer. These problems require different forms of review. Item-writing guidelines provide a vocabulary for identifying them (Haladyna et al., 2002; Tarrant et al., 2006), yet applying it across large item banks remains demanding.

Large language models (LLMs) change the scale of authoring without eliminating the need for validation (Balepur et al., 2025). Results in programming education demonstrate that GPT-4 can generate promising questions aligned with learning objectives (Doughty et al., 2024). The relevant next question is how to determine which generated or existing items are suitable for a particular assessment. Textual fluency, agreement with an answer key, and appropriate measurement of knowledge are distinct properties. Reference-overlap metrics such as BLEU and ROUGE measure similarity to reference text, not whether distractors are defensible or a question measures its intended construct (Moore et al., 2024). Human judgments also depend on the evaluation protocol and raters' expertise (van der Lee et al., 2021).

The Item-Writing Flaws (IWF) approach makes quality review more interpretable by assigning named flags to specific problems (Tarrant et al., 2006). The Scalable Automatic Question Usability Evaluation Toolkit (SAQUET) operationalizes a 19-criterion version of the IWF rubric using rules, NLP, and selective LLM verification (Moore et al., 2024). Related work extends automated review toward revision, psychometric and Item Response Theory (IRT) analysis, and auditing the MCQs used to evaluate language models. These applications share a rubric vocabulary, but make different evidentiary claims.

We contribute a criterion-level map of the automated application of IWFs, a comparison of detection and downstream evidence, and a research agenda connecting flags to assessment decisions. We distinguish *detecting a rubric violation*, *improving an item*, and *improving measurement*, noting that success at one stage does not establish success at the next. In measurement terms, a removed flaw is a change to the item, not yet evidence about the scores it produces.

2 Review Scope and Method

We conducted a focused narrative evidence map around five seed reports covering detection, remediation, psychometric and cognitive outcomes, and benchmark auditing. Searches of PubMed and Crossref on September 1, 2026 supplemented this set with records published from January 1, 2020 through that date. PubMed provided access to

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 464–471
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

biomedical and medical education literature on MCQ design and quality evaluation. Crossref provided complementary discovery through multidisciplinary bibliographic metadata, helping identify work beyond biomedicine (Gusenbauer, 2024). Targeted searches of education and computing venues and citation-based discovery completed the search. This bounded strategy does not establish exhaustive coverage of the relevant literature.
The title and abstract query was:

> ("multiple choice" OR "multiple-choice") AND ("item writing" OR "item-writing" OR "quality") AND ("large language model" OR "ChatGPT" OR "automated" OR "natural language processing" OR "NLP").

Crossref's `query.bibliographic` searches used *item-writing flaws* and *multiple choice question quality large language model*, retaining the first 25 relevance-ranked records per query. Exploratory searches targeting ACL Anthology, ACM, Springer, and ERIC, followed by title and DOI lookups, helped locate sources. References in the retrieved papers identified the earlier comparison by Moore et al. (2023).

Eligibility required an empirical evaluation of automated MCQ flaw detection or correction, or an analysis connecting automated flaw annotations to authoring, psychometric, or benchmark outcomes. Generation-only studies, model exam-taking studies without an evaluated flaw-screening component, general reviews, and unrelated assessment formats were excluded from the core synthesis. We included related technical-flaw rubrics without assuming equivalence to all 19 IWF criteria.

The databases returned 212 records, 162 from PubMed and 50 from Crossref. One duplicate DOI and one preprint/published-version pair left 210 reports for screening. Seven research papers met the criteria and 203 were excluded. Adding five seed reports, one citation-linked predecessor (Moore et al., 2023), and one subsequently nominated report (Arif et al., 2024) yielded fourteen included reports, all reviewed in full text (Table 1). Extraction distinguished the item sample, reference standard, rubric version, detector, metric denominator, and outcome. We prioritized criterion-level performance, positive class F1, and chance corrected agreement. Missing metrics remained unreported and aggregate accuracy was not interpreted as recall. Dataset reuse and incompatible units precluded summing sample sizes or pooling performance.

3 What the Rubric Operationalizes

IWF criteria describe avoidable sources of cueing, ambiguity, or irrelevant task demands (Tarrant et al., 2006). They do not constitute a complete theory of assessment quality. Content coverage, curricular alignment, fairness, and intended cognitive demand require additional evidence. An easy item may be appropriate for checking prerequisite knowledge, and a difficult item may be valuable for distinguishing advanced performance. Table 2 uses SAQUET's definitions as the reference vocabulary (Moore et al., 2024). Its routes illustrate three levels of automation. Text checks address visible patterns or format and NLP adds relationships among words and options. Selective LLM verification attempts judgments for which a pattern alone is insufficient. These routes describe implementation choices, not a fixed ordering of pedagogical importance. For example, absolute terms are easy to locate but may be entirely appropriate in a mathematical statement.

The rubric's lineage should not be mistaken for identical operationalization. Costello et al. (2018) describe 15 manually assessed flaws alongside programmatic item characteristics. SAQUET reports different human thresholds for the longest-correct-option criterion across its source datasets (Moore et al., 2024). BenchMarker separately represents a focused question and a problem fully stated in the stem, lacks SAQUET's explicit true/false criterion, and broadens the length rule to balanced options (Balepur et al., 2026). Several medical evaluations also adapt National Board of Medical Examiners (NBME) guidance (Sabqat et al., 2025). Accordingly, we preserve source definitions rather than treating equally named scores as interchangeable.

The newer applications expose further boundaries. Firoozi et al. (2025) implement six guidelines, including option-count limits and ordering options by character count, which differs from SAQUET's numerical or chronological sequencing rule. Arif et al. (2024) extend automated review to lecture transcripts with *contextual specificity*, where a question may depend on a diagram visible in the video but absent from the text. Such an item can work within a video yet fail as a standalone question. This demonstrates how rubric applicability, in part, depends on the information available to the learner.

Item format matters equally, such as the case when multiple defensible answers violate a single-best-answer item, but not an intentionally scored

Report / evidence	Sample and reference standard	What the evaluation establishes
Moore et al. (2023)	200 student-authored MCQs; two expert raters	Rules vs. GPT-4; domain micro-F1 .30–.70 vs. .25–.36.
Moore et al. (2024)	271 MCQs, five domains; existing expert labels	SAQUET: label accuracy 94.13%, exact match 38%; criterion evidence in Table 2.
Arif et al. (2024)	1,709 GPT-3.5 and 1,850 GPT-4 MCQs; 100 human-rated items	Rule/LLM rubric audits of video-based questions; missing context and distractor flaws; pass rates are not detector accuracy.
Sabqat et al. (2025)	13 NBME examples, one per flaw category	GPT-3.5 detection/correction before and after guideline prompting; no representative negative-item sample.
Aydin et al. (2026)	85 internal-medicine MCQs; three-expert reference	Three LLMs; negative-stem $\kappa = .897$ for Claude and Gemini; learning-outcome alignment ICC .121–.380.
Gérard et al. (2026)	Clinical pharmacology; 87 entries in the diagnostic 2x2 table	Shared model/key disagreement: sensitivity 35.3%, specificity 92.9%, precision 54.6%, NPV 85.5%.
Iembo et al. (2026)	36 original/revised pairs; 38 criteria; external and internal panels	External ratings found no significant aggregate change; calibrated author ratings favored revised items.
Menze et al. (2026)	152 items, 19 teachers; two LLMs; 30 human-rated items	AI/human-consensus Hamming loss .16; AI-supported items had more rubric flags.
Schmucker and Moore (2026)	7,126 STEM MCQs; estimated IRT parameters	Automated flags predict some parameter variation; difficulty R^2 reaches .216 overall but .030 in mathematics.
Moore (2026)	1,230 MCQs, three datasets; item-level response time and difficulty	Confusion-category flags associate with 14.0–29.4% longer response time; observational, category-level evidence.
Balepur et al. (2026)	12 NLP MCQ benchmarks; human annotation decisions	Rubric judges and benchmark filtering; NLP writing-flaw F1 .66, $\kappa = .53$ for Gemini 2.5 Pro.
Firoozi et al. (2025)	20 macroeconomics MCQs; two item-writing experts and one subject expert	Six-rule/GPT-4o revision: 17/20 flawless, 18/20 linguistically acceptable, 14/20 retain intended content.
Düsterbeck et al. (2026)	30 German biochemistry MCQs; six experts compare two refinements	Optimized agent sequences preferred for stems over a fixed pipeline; no option advantage.
Wu et al. (2026)	43 defective dental MCQs; 44 flaws identified by two experts	GPT-4 detects 12/44 flaws (27.3%); 7/7 duplicates but 1/23 off-topic flaws; no defect-free controls.

Table 1: Evidence map, not a system ranking, with all rows drawing on full text reports. Samples, rubric definitions, model versions, and outcomes differ. NPV is negative predictive value, ICC is intraclass correlation, and reported performance applies to the stated task and reference standard.

multi-select task. Similarly, the use of underscores in code should not automatically trigger a sentence completion flag. Quality assurance should record the item format and intended learner population before applying the rubric, and permit an explicit not applicable or uncertain judgment. Binary outputs without that context can convert legitimate design choices into apparent defects.

4 Evidence across the Assessment Workflow

4.1 Detection and the accuracy denominator

SAQUET combines deterministic checks with selective model judgments, while other systems prompt LLMs directly with an item-writing rubric or combine model judgments (Moore et al., 2023, 2024; Aydin et al., 2026; Menze et al., 2026). Both approaches are feasible, but comparisons across different datasets cannot establish architectural superiority or isolate benefits of newer models.

SAQUET makes the denominator problem concrete, as its 94.13% accuracy concerns 5,149 item-criterion decisions, not the fraction of flawed items successfully identified. The published human counts contain 356 positive labels. An always-negative classifier would therefore score $1 - 356/5149 = 93.09\%$ accuracy, a baseline calculated here. SAQUET's 38% exact match, domain micro-F1 of .59–.67, and 75.3% agreement on a separate acceptability threshold describe different tasks (Moore et al., 2024). Table 2 shows why averaging conceals useful distinctions.

Criterion	Operational concern	Route	Positive Label (n)	F1 range	IRT link
Longest option correct	Answer stands out through extra length/detail	T	59	.60–1.00	$b\downarrow, a\downarrow$
Ambiguous information	Wording permits unclear interpretations	N	40	.00–.58	—
Implausible distractors	Foils can be rejected without target knowledge	N	31	.00–.86	—
True/false format	Options reduce to binary truth judgments	T	27	.00–1.00	—
Absolute terms	Unqualified extremes provide a clue	V	25	.00–.89	$a\downarrow$ (Math)
Complex/K-type	Responses require combinations of statements	V	25	.00–.89	—
Negative wording	Negation complicates the requested judgment	T	23	.00–.91	$b\uparrow$
Convergence cues	Recombined option components suggest the key	V	19	.00–.63	$b\uparrow$
Lost sequence	Quantities/dates lack an intelligible order	T	19	.50–1.00	$a\uparrow$ / $\downarrow$
Unfocused stem	Stem does not adequately specify the task	V	16	.00–.89	—
None of the above	An exclusion option replaces a substantive answer	T	13	.91–1.00	—
Word repeats	Stem wording selectively points to the key	N	13	.53–1.00	—
Multiple correct answers	A single-best-answer item lacks a unique key	V	12	.00–.46	$b\downarrow, a\downarrow$
Logical cues	Option relations allow unintended deduction	N	8	.00–.67	—
All of the above	A collective option rewards partial recognition	T	6	.80–1.00	$b\downarrow, a\downarrow$
Fill in the blank	Mid-stem gaps interrupt the question	T	6	1.00	—
Vague terms	Unquantified expressions lack a shared meaning	T	6	.00–1.00	—
Grammatical cues	Only some options grammatically fit the stem	N	5	.00–.67	$b\uparrow$
Gratuitous information	Details add demands unrelated to the task	V	3	.00–.50	—

Table 2: Criterion map anchored in Moore et al. (2024). Routes: T = text/format checks, N = NLP, V = NLP with LLM verification. Positive labels (n) count items humans marked as having each flaw, out of 271. F1 ranges span domains (zeros included, uncomputed cells omitted; no pooling). IRT links from Schmucker and Moore (2026): b = difficulty, a = discrimination, mixed arrows = domain variation, and — = no link extracted, not evidence of no association.

Independent medical evidence reinforces the need for more than accuracy. Aydin et al. (2026) report strong negative-stem agreement for two models but much weaker alignment with learning objectives. In another approach, two models giving the same answer contrary to the key prompted a review of possible ambiguity. Its specificity was 92.9%, but sensitivity was only 35.3% (Gérard et al., 2026). Such a signal can focus attention while leaving many defects undetected. A model's failure to solve an item is not itself evidence that the item is defective.

Wu et al. (2026) quantify detection within a dental question-generation study. GPT-4 detected 12 of 44 expert-identified flaws across 43 defective items, consisting of 7/7 duplicates, 3/8 flawed distractors, 1/23 off-topic flaws, and 1/6 incorrect or missing answers. This 27.3% flaw-detection rate is not item-level accuracy, as without defect-free controls, specificity cannot be estimated. Despite requesting revisions, the study reports no separately validated repair-success rate.

Evaluations should report positive class support and per-criterion confusion matrices by domain, with precision, recall, specificity, F1, and uncertainty intervals. Chance corrected agreement is also sensitive to prevalence. Exact match and Hamming loss summarize label completeness and disagreement, but cannot replace criterion-level results. Deliberately flawed examples alone cannot establish

specificity on ordinary item banks (Sabqat et al., 2025).

In Arif et al. (2024), rules audit 19 criteria, where LLMs check seven IWFs plus contextual specificity on 100 items per model pool. Rules flag about 1.6 flaws per item, LLMs flag 2.18 for GPT-3.5, and 1.86 for GPT-4. These are automated flag counts, not error rates against expert labels. The human rubric also differs, so its 51-65% flawless item ratings cannot establish relative detector accuracy.

4.2 Revision and author interaction

Firoozi et al. (2025) combine six rule checks with chained GPT-4o revision and answer-key verification. Their 20-item macroeconomics pilot reports 17 flawless outputs, 18 linguistically acceptable outputs, and 14 preserving the intended content. Two item-writing experts evaluated quality, while a subject expert assessed content alignment. Ratings followed consensus and adjudication of the ordering rule. These are output-quality endpoints, not estimates of detector sensitivity or specificity. An essential point here is that a fluent repair may change what the question tests.

Düsterbeck et al. (2026) optimize criterion-specific agent sequences using REINFORCE algorithm and an LLM quality score. Six experts compared two refinements of 30 German biochemistry MCQs, initially generated by Qwen2.5-14B-Instruct to violate quality criteria. Optimized sequences were preferred for stems, but showed no advantage for options. The comparison is with a fixed refinement pipeline, as intentionally flawed inputs limit transfer to ordinary authoring and educator time savings were not measured.

Evidence also varies with the evaluators. In Iembo et al. (2026), ten external evaluators found no significant aggregate difference between original and revised questions, whereas calibrated ratings by four study authors favored the revisions. Calibration and evaluator group changed together, so the difference does not isolate a calibration effect. In Menze et al. (2026), AI-supported authoring was associated with more flags and less interaction. This work suggests fixed task order, self-selection into editing, and imperfect automated labels limit causal interpretation, including attribution to automation bias.

Based on this prior work, we believe studies should preserve original and revised items, distinguish accepted from rejected suggestions, and obtain blinded judgments of whether the intended defect was removed without introducing another. Content correctness, learning-objective alignment, and distractor function need separate checks. Counting generated revisions or repeat tool use does not establish successful remediation.

4.3 Psychometric and cognitive grounding

Schmucker and Moore (2026) connect automated annotations for 7,126 STEM MCQs with IRT parameters estimated from a platform serving 448,000 students. Associations vary by criterion and domain (Table 2). In the reported regression table, difficulty prediction reaches $R^2 = .216$ overall and only .030 in mathematics. A low-difficulty screening example reaches precision .801 with recall .472. These are useful trade-offs for triage, but leave substantial unexplained variation and missed cases.

The association between a flag and difficulty is not the effect of fixing that flag. Subject matter, reading demand, item purpose, and correlated flaws may explain part of the relationship. Moreover, the predictor itself is a fallible automated judgment. New-domain annotation checks are therefore needed before interpreting a regression coefficient as evidence about an underlying flaw. Concept specific IRT calibration also warrants attention when combining parameters from different content areas.

Moore (2026) extends this account by grouping SAQUET's 19 flags into cue-giving, confusion-inducing, and structural-complexity categories. The confusion category comprises ambiguity, unfocused stems, gratuitous information, lost sequence, negative wording, and vague terms. Across three datasets consisting of 653 USMLE, 365 XES3G5M mathematics, and 212 DBE-KT22 database-course items, its presence associates with 14.0%, 29.4%, and 18.9% longer mean response time, respectively, after structural controls. Cue-giving and structural indicators have no significant response-time associations. These are category-level estimates, not validated weights for each constituent criterion.

The difficulty pathway is more conditional. In USMLE, the reported confusion-category indirect estimate through response time is .231, approximately 2.5 times the aggregate estimate of .093. The corresponding category-specific indirect estimates are not significant in XES3G5M and DBE-KT22 uses instructor-assigned ordinal difficulty and a separate attenuation analysis (Moore, 2026). Thus, the 2.5-fold result is specific to USMLE. Item-level observational data cannot establish a

causal cognitive mechanism, and non-significant interactions do not demonstrate equivalent effects across domains. Machine-translated mathematics text, heterogeneous timing measures, and unvalidated new-domain flags further qualify transfer. The findings motivate targeted review, while controlled comparisons of original and repaired items are still needed to establish measurement benefits.

4.4 NLP benchmark auditing

BenchMarker extends educational review to twelve MCQ benchmarks, while separately checking online overlap and answer-option shortcuts (Balepur et al., 2026). Its reported tables show how accuracy can conceal weak positive-class performance and how results change between educational and NLP data. For example, Gemini 2.5 Pro obtains writing-flaw F1 of .66 on NLP annotations and .39 on educational annotations. These values concern annotation decisions rather than whole-item correctness.

The benchmark analysis also finds that filtering writing errors changes model rankings beyond random filtering. This makes item construction relevant to the interpretation of NLP evaluations. However, filtered subsets may differ in content and difficulty. Matched-size controls address sampling variation but do not establish a causal effect of correction. Paired repairs, preserved task coverage, and re-evaluation of all models would provide stronger evidence. Online occurrence should likewise be distinguished from demonstrated inclusion in a particular model's training data.

Balepur et al. (2026) also evaluated LLM-based repair of TruthfulQA items using detected item-writing flaws as feedback. Across three models, the proportion of items with at least two detected flaws decreased from 97% before revision to 64–68% afterward. However, BenchMarker identified newly introduced flaws in 56–68% of revised items. Because these outcomes were assessed automatically, they do not establish expert-confirmed repair quality or improved measurement. These findings motivate evaluating both the removal of existing defects and the introduction of new ones when assessing automated revision.

5 Research Priorities

Make definitions and uncertainty portable. Reusable datasets should release exact rubric wording, applicability rules, model versions, and positive and negative examples. Expert annotations should retain disagreements before adjudication and record the intended population and item purpose. Evaluation sets need enough positive cases for each criterion, including semantic defects. Table 2 shows that six criteria have fewer than ten positive cases in the SAQUET reference set, which limits what any F1 value for them can establish. Challenge sets should be reported separately from samples representing deployment prevalence. Shared items across dataset versions should have stable identifiers to prevent leakage and mistaken claims of independent replication.

Prioritize consequences, validate the policy. A raw flag count gives a rare grammatical cue and a second correct answer equal weight. Cognitive-category evidence motivates prioritization (Moore, 2026), but a review policy must also consider whether a flag is correct, its consequence for the intended assessment, and the cost of checking it. The category results do not validate universal severity weights or imply that cue-giving defects are harmless. Evaluate priorities on held-out domains and independently judged revisions, with student response evidence where appropriate. LLM confidence should not be assumed to be a calibrated probability.

Evaluate the complete human-AI workflow. Compare unaided review, flags alone, flags with explanations, and suggested revisions. Measure expert-confirmed defect removal, introduced defects, reviewer time, and escalation. Expert review of unflagged items should estimate missed problems. Subsequent pilots can examine discrimination, intended difficulty, response time, and subgroup performance, connecting authoring efficiency to measurement outcomes.

Design paired item comparisons. Paired evaluations should preserve the learning objective while changing the targeted flaw and assign learners to one version only. Shared anchor items can support comparisons on a common scale. Removing a cue may increase difficulty while improving validity, so neither faster responses nor easier items should define success in isolation. Record rejected suggestions and reasons alongside accepted revisions; acceptance alone can reflect misplaced trust.

6 Limitations

This review is limited by bounded Crossref retrieval, database coverage, nominated-source selection, and AI-assisted screening without independent replication. The corpus is not exhaustive, and the first author co-authored five of the fourteen included reports, including the toolkit that anchors Table 2. We did not rerun detectors or reanalyze participant data. Much of the evidence uses small or reused annotation sets, and comparisons span different rubric versions and evaluation units. The cognitive findings are observational and the proposed workflow and consequence-based priorities still require validation.

7 Conclusion

AI-enabled IWF review offers interpretable support for screening MCQs, but its evidence is uneven across criteria and stages of use. Surface checks are often tractable, but content sensitive judgments and reliable revision remain harder. The strongest path forward is to evaluate detection, remediation, and assessment outcomes separately and then test how they work together. Rubric flags can organize human attention before deployment, while expert review and response-based validation establish whether the resulting items serve their intended purpose.

References

Taimoor Arif, Sumit Asthana, and Kevyn Collins-Thompson. 2024. Generation and assessment of multiple-choice questions from video transcripts using large language models. In *Proceedings of the Eleventh ACM Conference on Learning @ Scale*, pages 530–534. Association for Computing Machinery.

Mevlüt Okan Aydin, Belkıs Nihan Coşkun, and İbrahim Hamal. 2026. Assessing multiple-choice question quality in internal medicine: A comparative analysis of three large language models against expert consensus. *Frontiers in Medicine*, 13:1866674.

Nishant Balepur, Bhavya Rajasekaran, Hyunjin Jane Oh, Michael Xie, Atrey Desai, Vipul Gupta, Steven James Moore, Eunsol Choi, Rachel Rudinger, and Jordan Lee Boyd-Graber. 2026. Benchmarker: An education-inspired toolkit for highlighting flaws in multiple-choice benchmarks. In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15793–15824.

Nishant Balepur, Rachel Rudinger, and Jordan Lee Boyd-Graber. 2025. Which of these best describes multiple choice evaluation with llms? a) forced b) flawed c) fixable d) all of the above. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3394–3418.

Eamon Costello, Jane C. Holland, and Colette Kirwan. 2018. Evaluation of MCQs from MOOCs for common item writing flaws. *BMC Research Notes*, 11:849.

Jacob Doughty, Zipiao Wan, Anishka Bompelli, Jubahed Qayum, Taozhi Wang, Juran Zhang, Yujia Zheng, Aidan Doyle, Pragnya Sridhar, Arav Agarwal, Christopher Bogart, Eric Keylor, Can Kultur, Jaromir Savelka, and Majd Sakr. 2024. A comparative study of AI-generated (GPT-4) and human-crafted MCQs in programming education. In *Proceedings of the 26th Australasian Computing Education Conference*, pages 114–123. Association for Computing Machinery.

Lilly Marie Düsterbeck, Michael Grössler, Graziella Credidio, Louis Bellmann, and Layla Tabea Riemann. 2026. Agent-based improvement of multiple-choice question quality in medical education. *Studies in Health Technology and Informatics*, 336:2120–2124.

Tahereh Firoozi, Lia Daniels, Vijay Daniels, and Mark Gierl. 2025. An augmented intelligence system for automated quality control and feedback generation of multiple choice test items. In *Artificial Intelligence in Education. Posters and Late Breaking Results, Workshops and Tutorials, Industry and Innovation Tracks, Practitioners, Doctoral Consortium, Blue Sky, and WideAIED*, volume 2592 of *Communications in Computer and Information Science*, pages 86–93, Cham. Springer.

Alexandre O. Gérard, Diane Merino, Marc Labriffe, Fanny Rocher, Delphine Viard, Laurence Zemori, Thibaud Lavrut, Erik M. Donker, Joost D. Piët, Jean-Paul Fournier, Milou-Daniel Drici, and Alexandre Destere. 2026. Evaluating and leveraging large language models in clinical pharmacology and therapeutics assessment: From exam takers to exam shapers. *British Journal of Clinical Pharmacology*, 92(8):2574–2583.

Michael Gottlieb, John Bailitz, Megan Fix, Eric Shappell, and Mary Jo Wagner. 2023. Educator's blueprint: A how-to guide for developing high-quality multiple-choice questions. *AEM Education and Training*, 7(1):e10836.

Michael Gusenbauer. 2024. Beyond google scholar, scopus, and web of science: An evaluation of the backward and forward citation coverage of 59 databases' citation indices. *Research synthesis methods*, 15(5):802–817.

Thomas M. Haladyna, Steven M. Downing, and Michael C. Rodriguez. 2002. A review of multiple-choice item-writing guidelines for classroom assessment. *Applied Measurement in Education*, 15(3):309–333.

Tatiane Iembo, Helena Landim Gonçalves Cristóvão, Patrícia Carla Zanelatto Gonçalves, Wagner Ricardo Montor, Patrícia da Silva Fucuta, Toufic Anbar Neto, Júlio César André, and Milton de Arruda Martins. 2026. ChatGPT as a tool for reviewing multiple-choice questions in the health sector. *Scientific Reports*, 16:21974.

Dennis Menze, Slaviša Radović, and Niels Seidel. 2026. AI-assisted MCQ creation increases item-writing flaws through automation bias: Evaluation of a workflow of multiple AI agents with teachers. *Frontiers in Computer Science*, 8:1831250.

Steven Moore. 2026. The cognitive cost of poor item design: How item-writing flaws shape response time and difficulty across assessment domains. In *Proceedings of the 19th International Conference on Educational Data Mining*, pages 293–305, Seoul, Republic of Korea. International Educational Data Mining Society.

Steven Moore, Eamon Costello, Huy A. Nguyen, and John Stamper. 2024. An automatic question usability evaluation toolkit. In *Artificial Intelligence in Education*, pages 31–46, Cham. Springer.

Steven Moore, Huy A. Nguyen, Tianying Chen, and John Stamper. 2023. Assessing the quality of multiple-choice questions using GPT-4 and rule-based methods. In *Responsive and Sustainable Educational Futures*, pages 229–245, Cham. Springer.

Mashaal Sabqat, Rehan Ahmed Khan, Masood Jawaid, and Madiha Sajjad. 2025. Artificial intelligence meets item analysis (AI meets IA): A study of chatbot training and performance in detecting and correcting MCQ flaws. *Pakistan Journal of Medical Sciences*, 41(3):652–656.

Robin Schmucker and Steven Moore. 2026. The impact of item-writing flaws on difficulty and discrimination in item response theory. *Computers and Education: Artificial Intelligence*, page 100632.

M. Tarrant, A. Knierim, S. K. Hayes, and J. Ware. 2006. The frequency of item writing flaws in multiple-choice questions used in high stakes nursing assessments. *Nurse Education Today*, 26(8):662–671.

Chris van der Lee, Albert Gatt, Emiel van Miltenburg, and Emiel Krahmer. 2021. Human evaluation of automatically generated text: Current trends and best practice guidelines. *Computer Speech & Language*, 67:101151.

Jiashun Wu, Huanzhong Ji, Yanqing Xie, Shijun Kuang, and Tao Yang. 2026. Evaluating the performance of GPT-4.0 in generating and refining multiple-choice questions for dental education. *European Journal of Dental Education*.

More than Words: Modeling What Teachers Praise and How They Praise It

Jiseung Yoo Paiheng Xu Jing Liu
University of Maryland, College Park, MD 20742, USA
{jyoo20, paiheng, jliu28}@umd.edu

Abstract

Praise is one of the most frequent evaluative acts in teaching, but what a praise turn communicates depends less on its words than on how it is delivered. Text-based measures cannot see this channel, and absolute acoustic measures confound stable teacher style with event-level delivery. An expressive teacher sounds animated during routine instruction, while a reserved teacher marks praise with a small change in pitch or timing. We propose measuring delivery as *self-deviation*, comparing a praise turn against two references drawn from the speaker themselves, the teacher's other praise turns (Global Reference) and the speech immediately surrounding the turn (Local Reference). We instantiate this in SERMA (Self-Reference Modeling of Affective Delivery), a low-resource model built over frozen SSL speech features.

1 Introduction

Teachers praise constantly. "Good job," "Nice work," and "Great thinking" recur throughout classroom and tutoring transcripts. Counting these expressions is straightforward, but determining what they communicate is not. A single praising utterance does several things at once. It carries information about what the teacher valued, evaluates the student against some standard of worth, and conveys the teacher's affect (Brophy, 1981; Brummelman et al., 2016; Stout, 2020). Because these functions travel together, the meaning of praise is not fixed in its words. It is a situated act whose effect depends on timing, on how the student perceives it, and above all on how the teacher delivers it (Morton et al., 2022; Wolff and Brechmann, 2015). Praising effort rather than results is generally more motivating, yet older students can read the same praise as a signal of low ability (Amemiya and Wang, 2018).

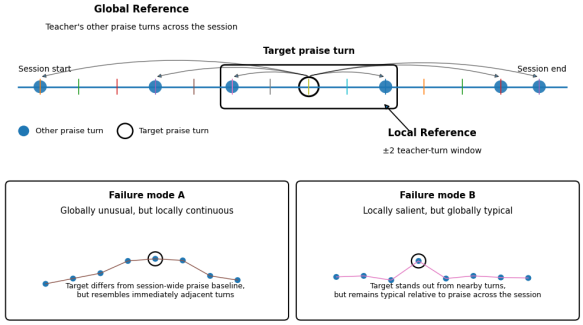


Figure 1: Conceptual scheme of global and local references.

Textually identical utterances differ sharply in felt authenticity (Burnett, 2002; Weinstein et al., 2020; Paulmann and Weinstein, 2025).

Measurement has lagged behind this understanding. Early work treated praise as reinforcement and counted frequencies (Becker et al., 1967). Attribution research turned to what is praised (Mueller and Dweck, 1998), and work on behavior-specific praise showed that naming the recognized action communicates more than general approval (Polick et al., 2012; Floress et al., 2018), while self-determination theory made quality rather than quantity the central concern (Ryan and Deci, 2000; Soenens and Vansteenkiste, 2020). The empirical literature nonetheless remains largely quantitative, counting praise events or coding wording as general versus behavior-specific (Chen et al., 2025). Text-based NLP can identify praise phrases at scale, but transcripts erase the acoustic channel through which teachers mark significance, warmth, and immediacy, and existing systems cannot distinguish a "good job" delivered warmly from the same words spoken flatly (Chen et al., 2025; Göllner et al., 2025).

The channel is known to be informative. Prosody carries warmth, energy, emphasis, and restraint

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 472–482
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

(Paulmann and Weinstein, 2025), and vocal information improves both human judgments and model performance beyond text alone (Onishi et al., 2020). What is missing is a way to turn that signal into a measure, because there is no universal acoustic definition of good praise. Autonomy-supportive delivery may be calm and restrained while other positive engagement is energetic. Absolute acoustic values therefore confound stable teacher style with event-specific delivery, which is precisely the quantity of interest.

Speech emotion recognition (SER) offers a route to separating the two. Norm theory holds that judgment compares an observed event against an available standard (Kahneman and Miller, 1986), and speaker-sensitive SER systems evaluate a turn against the speaker's own distribution rather than a population threshold. The gain comes from what normalization removes, since a pitch or intensity value that signals arousal in one voice falls within the resting range of another, and only the residual after conditioning on the speaker is comparable across speakers. Conversational models extend the same logic to surrounding talk, showing that speaker identity and local context jointly determine how a turn is heard (Shi et al., 2025; Du et al., 2025; Zhang and Tan, 2025; Jeong and Nyamathi, 2026). Together these imply that enthusiastic delivery is not an absolute property of a voice but an event-level deviation from the speaker's own pattern against the surrounding talk (Martinez-Lucas et al., 2026; Tamir et al., 2025).

Speaker normalization is not itself new (Gat et al., 2022; Shi et al., 2025). What the construct contributes is a principled answer to a question that normalization alone leaves open: what to normalize against. We argue that the reference set must be selected by the construct rather than by convenience. Because the question is whether this praise stands out among the ways this teacher normally praises, the Global Reference is the teacher's other praise turns rather than their speech in general, which would fold the ordinary prosodic contrast between instruction and evaluation into the measure and report it as enthusiasm; work on conversational annotation likewise conditions judgments on the speaker's own characteristics. The Local Reference addresses a complementary failure mode. A turn can be globally unusual yet locally continuous, as when the teacher has been animated for the past minute, or locally salient yet unremarkable for a routinely expressive teacher; speakers align their prosody to the immediate talk rather than to a fixed personal baseline, so the surrounding discourse is the operative norm at that moment (Ostrand and Chodroff, 2021). This two-axis, construct-selected design is what separates the measure from a generic emotion classifier, which registers whether a voice is loud or high rather than whether a turn departs from the speaker's own pattern.

Our contributions are threefold. (i) We present a measurement framework that operationalizes praise delivery through theory-informed reference-relative representations. Theory fixes the construct and its reference relation before any model is fit, the representation encodes that relation, and predictive accuracy is then used to test the representation rather than to select it. (ii) Because praise is both what is said and how it is said, we develop annotation rubrics and an identification pipeline for both the linguistic and the acoustic channel, and we use the linguistic channel to define the praise-relative reference set. Modeling the two channels jointly is the natural next step and is not attempted here. (iii) We develop SERMA, a low-resource method combining frozen self-supervised speech features, two reference channels, teacher-grouped evaluation, and strategies for limited annotation and severe class imbalance (Ueda et al., 2026; Naini et al., 2025; Gong et al., 2025). Freezing the backbone reduces overfitting and annotation cost, which matters because classroom corpora are hard to annotate; layer-wise analyses of frozen SSL features show that mid-layer prosodic content with statistical pooling, retaining the variance, carries the within-turn variability that enthusiasm lives in (Maji et al., 2024; Uniyal and Abrol, 2025), and naturalistic benchmarks establish that imbalance-aware objectives and uncertainty-aware label handling are needed when positive events are scarce (Omidi and Hansen, 2026).

This paper reports an in-progress instrument. We present the construct, the annotation protocol, and a first evaluation of reference-relative representations against absolute-feature and zero-shot baselines under leave-one-teacher-out cross-validation. The ablation that zeroes both deviation channels, and the paired cluster-bootstrap intervals over per-fold differences, are specified below but not yet run, so the results here establish an ordering across a family of representations rather than a causal account of where the gain originates.

2 Data and Annotation

2.1 Data

We analyze 157 elementary tutoring sessions from five teachers in an online early-literacy program. All taught Grades 2 to 3 in groups of two or three. Sessions averaged around 40 minutes. Audio came from Zoom recordings with speaker-attributed, timestamped human transcripts, yielding 21,843 teacher turns from a total of 46,437 turns. The setting suits measurement development because it offers sustained interaction, repeated evaluative exchanges, and many observations per teacher. Praise is frequent. Teachers produced 0.84 praise turns per minute, and 24.7% of teacher turns contained a praise expression.

2.2 Annotation

We annotate two constructs, one over the words of a praise turn and one over its delivery. Table 1 summarizes both.

For **linguistic specificity**, four coders worked over a teacher-stratified random sample of 800 candidate turns. They first judged whether each candidate functioned as praise, and the turns judged to be praise were then coded for specificity. A turn counts as specific if it names an action or effort (C1), echoes the student's contribution (C2), or personalizes recognition by name or pronoun (C3). Meeting any one criterion sufficed. Coders saw three preceding turns and one following turn. First-round agreement was Fleiss' kappa 0.803 for the praise decision and 0.743 for specificity, and rose to 0.854 and 0.797 after a rubric revision round. Per-criterion agreement for both rounds appears in Table A1.

For **acoustic enthusiasm**, three coders rated 757 praise turns on a five-point scale anchored at no enthusiasm (1) and high enthusiasm (5), with the neutral midpoint (3) removed, leaving four options, so that coders committed to the lower or upper half of the range. Coders heard audio with synchronized transcripts. Turns were batched by teacher and session, so each coder heard a teacher's full set before rating any single turn. First-stage agreement was moderate (Fleiss' kappa 0.471). After a teacher-grouped pairwise calibration stage on 201 difficult turns, agreement over the full set reached 0.629. The final annotation gave 113 positive and 644 negative turns, a positive rate of 14.9%.

2.3 Label–construct alignment, and what it costs

Since coders heard a teacher's full set before rating any single turn, the labels are themselves speaker-relative, which could advantage a speaker-relative model by construction. We treat this as a consequence of the construct: if enthusiasm is defined relative to a speaker, a protocol that induces that standard is the right operationalization, whereas shuffling turns across teachers would elicit judgments of absolute vocal arousal. The cost is that the present results cannot separate how much of SERMA's advantage comes from the construct and how much from the protocol that produced the labels. We plan to re-rate a subset of turns under a shuffled protocol to estimate how far judgments move when the speaker-relative standard is removed.

3 Method

3.1 Linguistic pipeline

A two-stage LLM pipeline, prompted with the rubric and few-shot examples, identifies praise and its specificity. For each target utterance we build a window of three preceding turns, the target, and one following turn, matching the coders' context. Stage 1 scans for evaluative language and extracts candidate phrases with temporal boundaries. Stage 2 decides whether each candidate functions as praise, then labels validated praise as specific or general by C1 to C3. Table 2 reports accuracy against the 800 human-coded candidate turns, and all downstream experiments use the praise decisions of the selected system there. Because these decisions are made without reference to enthusiasm, they do not leak the target construct.

3.2 Acoustic pipeline

3.2.1 Utterance encoder

Each turn is encoded from frozen SSL speech frames, WavLM layers 0 to 8, $X_i \in \mathbb{R}^{T_i \times D}$ ($D = 1024$). We restrict to lower layers because they retain more prosodic and less lexical information, and restricting rather than letting the mixture discover this is a regularization choice given 757 labels. A learned softmax over per-layer logits mixes retained layers, so layer selection is optimized with the downstream model while SSL weights stay fixed. Frames are projected to $h = 64$ and passed to a one-layer unidirectional GRU, conditioning

Table 1: Annotation taxonomy.

Channel	Component	What it marks	Example
Linguistic (what is praised)	Action verb (C1)	Names the specific action or effort, supporting autonomy by making the recognized behavior explicit	"You solved the problem perfectly" (Kamana et al., 2022; Stevens et al., 2011)
	Echo (C2)	Repeats or rephrases the student's contribution, supporting competence by reflecting the work back	Teacher restates the student's answer with affirmation (Brophy, 1981; Granger et al., 2024)
	Personal (C3)	Names the student or addresses them directly, supporting relatedness	"Great job explaining that, Sarah" (Kamana et al., 2022; Rila et al., 2024)
Acoustic (how it is delivered)	Expressive delivery	Variation in pitch, volume, and tempo, with rise–fall contours and sustained intensity at the peak	A wider pitch range than the speaker's baseline, roughly 72 to 147 Hz (Weyman and Sy, 2018; Wolff and Brechmann, 2015)
	Contrast (global and local)	Departure from how this teacher usually praises and from the speech immediately surrounding the turn, giving selective signaling and the perceptual pop effect	Moderate to high intensity shift against the speaker's own praise baseline and against the adjacent turns (Paulmann and Weinstein, 2025)

each frame on preceding acoustics. This frame-level module and the turn-level reference encoder described below are the two components varied across the model family, and SERMA-Transformer replaces both with transformer encoders. The hidden sequence is pooled by a masked-softmax scorer that weights valid frames, and the weighted mean and standard deviation are concatenated and projected to $z_i \in \mathbb{R}^h$. Retaining the second moment preserves within-turn variability, which mean pooling discards and which carries much of the prosodic dynamics of enthusiasm. Crucially, z_i is purely acoustic, and no lexical information enters, so the deviations below compare prosody against prosodic baselines rather than semantic drift. Reference turns are encoded by this same frame encoder with shared weights, so reference embeddings occupy the same space as the target. Sequence-length handling, dropout, and remaining dimensions are in Appendix B.

3.2.2 Reference sets

The two reference sets use deliberately different scopes because they answer different questions. The **Global Reference** set comprises the teacher's own other praise turns, uniformly subsampled to at most $K_g = 24$ and excluding the target. The construct asks whether this praise stands out among the ways this teacher normally praises, so the baseline must be teacher praise rather than teacher speech in general. The **Local Reference** set is a symmetric window of $K_l = 4$ teacher turns, the two preceding and the two following, capturing whether the turn stands out against surrounding speech at that moment. Both sizes were fixed before evaluation and never tuned on a held-out teacher. Neither set uses enthusiasm labels. Local membership derives from timestamps and speaker metadata alone, and global membership additionally uses the automatic praise decision, which is made independently of the target construct.

3.2.3 Deviations and availability gating

Reference embeddings are averaged into z_l and z_g, and the deviations $d_l = z - z_l$ and $d_g = z - z_g$ are each encoded as $[d,\ |d|]$ by a linear-GELU layer. The signed and absolute parts are both retained because direction and magnitude of departure are separately informative. Encoded deviations are then rescaled by availability gates computed from a descriptor of reference-set existence and size (Appendix B), so a deviation resting on few references is attenuated rather than consumed as a fabricated zero. Gating is load-bearing here, because a praise-filtered global set is sparse far more often than one drawn from all of a teacher's speech. Writing l and g for the gated encoded local and global deviations, the two deviations recombine with the target

through a residual projection $c = z + W\,[z,\ l,\ g]$.

3.2.4 Context, fusion, and the two heads

A turn-level encoder, a GRU in SERMA and a transformer encoder in SERMA-Transformer, reads the ordered Global Reference sequence terminated by the target, and its output at the target position is projected and added to c. This encoder supplies how the teacher's praise delivery drifts across a session, rather than local discourse context, which the local deviation already carries. Where transcripts exist, a linear projection of a 774-d per-turn RoBERTa embedding is added at this same point, after the acoustic deviations, keeping the reference pathway audio-only and the lexical contribution separable, with the consequence that the target head stays audio-only even in the audio+text condition. Every text and audio+text row reported below uses this same text encoder. Two linear heads then train jointly, a **target head** reading z alone and a **context head** reading c. The loss is the mean of the two (BCE, or class-weighted cross-entropy with inverse-frequency weights fit on the training fold), and at inference the two head probabilities are averaged to give the score that is calibrated and thresholded as described below. The target head is not redundant. Keeping a path from acoustics to prediction that carries no reference information means the model is never *forced* to define enthusiasm as unusualness for this teacher, and if absolute delivery is what annotators responded to, that path can express it.

4 Experimental Setup

4.1 Protocol

We use leave-one-teacher-out cross-validation. Each fold trains on four teachers and tests on the fifth, so no teacher appears in both splits. This matches the construct, since a model measuring deviation from a teacher's own baseline should transfer to a new teacher with no labeled data from them. Within each fold, 20% of training turns (approximately 120) form a calibration set used only for calibration and threshold selection. The systems differ in how a score is produced. For the fixed-feature representations (Praat, WavLM-mean, FixedRef) the classifier is L2-regularized logistic regression (balanced weights, $C = 1.0$, in-fold standardization), chosen for stability with high-dimensional features and few labels. The SERMA variants are trained end-to-end with the two prediction heads and joint loss described above, and their averaged head probability is the score. Given the 14.9% positive rate, we apply sigmoid (Platt) calibration, since isotonic overfits a calibration set this small; we then sweep the threshold from 0.1 to 0.9 in steps of 0.05 and apply the best value once to the held-out teacher. The zero-shot scores are thresholded on the same calibration folds under this same procedure, using no held-out labels. With roughly 18 positive turns in a calibration set, this sweep is itself a source of fold-to-fold variance, and a fixed rule that sets the threshold at the training base rate is an alternative we have not yet compared against.

4.2 Compared systems

Praat uses hand-crafted descriptors (Appendix B). *WavLM-mean* uses mean-pooled frozen WavLM frames with no sequence model. *FixedRef* implements the same two reference relations in a fixed basis. Frames are projected by an unsupervised PCA to 64 dimensions, global and local deviations are taken in that fixed space and retained at full dimensionality, KL-divergence statistics over the two reference distributions and their interaction are appended alongside prosody-shape descriptors and an absolute content channel, and the resulting stack is classified by the same logistic regression used for the other fixed-feature representations. FixedRef is therefore the control for how the relations are encoded rather than for whether they matter at all, which is the role of the remove-both ablation. *SERMA-Transformer* replaces both recurrent modules, the frame-level encoder and the turn-level reference encoder, with transformer encoders (Appendix B).

Three zero-shot systems run on the same held-out teachers. **ZS-Point** prompts GPT-5.4 to score a single turn from Praat z-scores rendered as prosody descriptions. **ZS-Pair-Audio** gives Gemini-3.5-flash-lite two praise turns from the same teacher, with audio and transcripts, and asks which is more enthusiastic, aggregating preferences by regularized Bradley–Terry. **ZS-Pair-Praat** repeats that protocol with Praat z-scores as text in place of audio. ZS-Pair-Audio versus ZS-Pair-Praat isolates audio perception from reasoning within one model. ZS-Point versus ZS-Pair-Praat compares pointwise with pairwise judgment, though these differ in vendor as well, so the comparison is not fully isolated.

4.3 Metrics

The primary metric is positive-class F1 at the selected threshold, with PR-AUC as a threshold-independent complement. Macro-F1 appears for comparability with prior work but never drives threshold selection, since negative-class F1 is high by default at this base rate and inflates the average even when enthusiastic praise is missed. Metrics are computed per held-out teacher and averaged across folds rather than pooled, so each teacher counts equally. All systems see identical folds, which permits paired per-fold differences with cluster-bootstrap intervals over teachers. We report point estimates here and leave the paired intervals to the next revision, since marginal intervals on five folds are wide enough to be uninformative about contrasts.

5 Results

5.1 Linguistic pipeline result

All three models detect praise at comparable F1 (0.878 to 0.887), so detection is not where they separate. They differ on specificity, where GPT-5.6-luna reaches 0.825 F1 and is the most even across the three criteria. Qwen-3.7-28B is close overall (0.816) but weaker on naming an action (C1, 0.707 against 0.754), while GPT-5.4 loses most of its accuracy on both C1 (0.622) and personalized recognition (C3, 0.628). We therefore use GPT-5.6-luna's praise decisions throughout, and note that it also has the highest detection precision (0.938 against 0.918 and 0.888), which is the quantity that governs how much non-praise speech contaminates the Global Reference set. Full results are in Table 2.

5.2 Acoustic pipeline result

SERMA is the strongest system at 0.542 F1 and 0.729 macro-F1 (Table 3). The ordering across the supervised family is stable and interpretable. Hand-crafted descriptors and mean pooling fail (0.303, 0.353), representations that encode the two reference relations reach 0.475 and above, and encoding those relations in a learned, gated space rather than a fixed PCA basis adds a further 0.067. Because FixedRef computes the same global and local deviations, this margin is attributable to how the relations are encoded rather than to their presence, and the complementary question of what the relations contribute at all is left to two ablations that are not yet run, one zeroing both deviation channels and one replacing the availability gates with a constant. One caveat governs how all of these margins should be read. The fold-to-fold standard deviations, 0.144 for SERMA and 0.180 for FixedRef, are of the same order as the differences between systems, so the three-tier ordering is consistent across folds while individual pairwise margins remain provisional.

Two further observations. Text alone is weak everywhere (0.272 to 0.329), which is the premise of the paper, since what was said is nearly uninformative about how enthusiastically it was said, though SERMA's text-only PR-AUC of 0.402 sits close to its audio-only 0.446, which suggests that much of the text-only collapse is thresholding fragility on roughly 18 calibration positives rather than an absence of signal. Audio alone recovers most of the combined performance, trailing the audio+text rows by 0.06 to 0.11 among the reference-relative systems and by at most 0.03 for the two absolute-feature systems. The zero-shot comparison is also instructive about where the difficulty lies. ZS-Pair-Audio at 0.471 matches the strongest supervised baseline without task-specific training, while ZS-Pair-Praat, which gives the same model the same protocol over textual z-score descriptions, collapses to 0.281. Listening, not reasoning, is what carries the judgment, and the pairwise framing is itself doing work, since it forces the model into the same within-teacher comparison that our reference relations encode.

6 Discussion and Future Directions

This study builds and evaluates an instrument for measuring praise delivery. We ask whether a theory-specified representation aligns with human judgments, not whether it maximizes prediction accuracy, and with five teachers the results show feasibility rather than generalizability.

Three results stand out. Representations that encode self-deviation reach 0.475 positive-class F1 and above, against 0.303 for hand-crafted prosodic descriptors and 0.353 for mean-pooled WavLM, and the ordering holds on all five held-out teachers. SERMA is the strongest at 0.542, but its 0.067 margin over the same relations computed in a fixed unsupervised basis is smaller than the fold-to-fold standard deviations of 0.144 and 0.180, and it differs from FixedRef in several ways at once, so the result establishes an ordering while the size of the gain remains open. The signal is also acoustic: audio alone recovers most of the audio+text result,

Table 2: Performance of praise detection and specificity models.

Model	Task	F1	Precision	Recall
GPT-5.4	Detection	0.878	0.888	**0.868**
	Specific	0.779	0.723	0.844
	Verb (C1)	0.622	0.491	0.849
	Echo (C2)	0.740	0.675	0.818
	Personal (C3)	0.628	0.696	0.571
Qwen-3.7-28B	Detection	**0.887**	0.918	0.859
	Specific	0.816	0.710	**0.959**
	Verb (C1)	0.707	0.592	**0.879**
	Echo (C2)	0.755	0.638	**0.923**
	Personal (C3)	0.746	0.759	0.733
GPT-5.6-luna	Detection	0.880	**0.938**	0.829
	Specific	**0.825**	**0.773**	0.885
	Verb (C1)	**0.754**	**0.722**	0.788
	Echo (C2)	**0.763**	**0.726**	0.803
	Personal (C3)	**0.764**	**0.778**	**0.750**

Table 3: Performance of enthusiastic praise detection models.

System	Modality	F1	±	Macro-F1	PR-AUC
Praat	audio+text	0.303	0.030	0.435	0.257
	audio	0.273	0.081	0.486	0.250
	text	0.285	0.058	0.468	0.308
WavLM-mean	audio+text	0.353	0.057	0.561	0.267
	audio	0.340	0.084	0.560	0.329
	text	0.272	0.051	0.410	0.217
FixedRef	audio+text	0.475	0.180	0.640	0.430
	audio	0.367	0.099	0.590	0.381
	text	0.306	0.128	0.427	0.218
SERMA	audio+text	**0.542**	0.144	**0.729**	**0.460**
	audio	0.481	0.157	0.686	0.446
	text	0.324	0.119	0.469	0.402
SERMA-Transformer	audio+text	0.519	0.152	0.689	0.433
	audio	0.453	0.197	0.630	0.402
	text	0.329	0.098	0.527	0.328
ZS-Point	Praat-as-text	0.352	n/a	0.541	0.244
ZS-Pair-Audio	audio+text	0.471	n/a	0.661	0.368
ZS-Pair-Praat	Praat-as-text	0.281	n/a	0.327	0.201

text alone is weak in every system, and a zero-shot model given audio reaches 0.471 against 0.281 for the same model given the same pairwise protocol over Praat z-scores as text. None of these comparisons isolates the reference relations, since every system that encodes them also differs in how it encodes them.

We plan four next steps. An ablation that zeroes both deviation channels, with the frame encoder, the turn-level encoder, and the prediction heads held fixed, would separate what the reference relations contribute from what the encoder carrying them contributes; a companion ablation replaces the availability gates with a constant to test whether attenuating deviations from sparse reference sets matters. A sensitivity check varies the global and local set sizes and substitutes another teacher's praise for the speaker's own, testing whether the measure tracks the baseline it is given. Paired per-fold differences with cluster-bootstrap intervals over teachers would remove the shared fold difficulty behind most of the spread in Table 3 and support comparisons between systems. Re-rating a subset of turns under a teacher-shuffled protocol (Section 2.3) would estimate how far judgments move once the speaker-relative standard is removed. More teachers are needed before any gain can be sized, joint modeling of specificity and delivery would test an interaction the construct predicts, and single-pass inference would make the

method usable for feedback during instruction.

References

Jamie Amemiya and Ming-Te Wang. 2018. Why effort praise can backfire in adolescence. *Child Development Perspectives*, 12(3):199–203.

Wesley C. Becker, Madsen, Charles H., Jr., Carole R. Arnold, and Don R. Thomas. 1967. The contingent use of teacher attention and praise in reducing classroom behavior problems. *The Journal of Special Education*, 1(3):287–307.

Jere Brophy. 1981. Teacher praise: A functional analysis. *Review of Educational Research*, 51(1):5–32.

Eddie Brummelman, Jennifer Crocker, and Brad J. Bushman. 2016. The praise paradox: When and why praise backfires in children with low self-esteem. *Child Development Perspectives*, 10(2):111–115.

Paul C. Burnett. 2002. Teacher praise and feedback and students' perceptions of the classroom environment. *Educational Psychology*, 22(1):5–16.

Zhongpu Chen, Decheng Ren, Zhongying Wang, Liyue Gao, Ying Shi, Haiying Liu, and Tao Yang. 2025. Enhancing student academic performance in middle school classrooms by fostering engagement motivation through intelligent assessment of teacher praise emotional intensity. *Learning and Instruction*, 99:102185.

Hao Du, Shiyin Lu, Guangyu Zhou, and Ye Gao. 2025. EAA: Emotion-aware audio large language models with dual cross-attention and context-aware instruction tuning. In *Proceedings of Interspeech 2025*, pages 5433–5437. ISCA.

Margaret T. Floress, Lyndsay N. Jenkins, Wendy M. Reinke, and Laura McKown. 2018. General education teachers' natural rates of praise: A preliminary investigation. *Behavioral Disorders*, 43(4):411–422.

Itai Gat, Hagai Aronowitz, Weizhong Zhu, Edmilson Morais, and Ron Hoory. 2022. Speaker normalization for self-supervised speech emotion recognition. In *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7342–7346. IEEE.

Richard Göllner, Rebecca Lazarides, and Patrick Stark. 2025. Revealing teaching quality through lesson semantics: A GPT-assisted analysis of transcripts. *British Journal of Educational Psychology*, 95:S300–S315.

Z. Gong, P. Shi, K. Donbekci, L. Ai, R. Chen, D. Sasu, Z. Wu, and J. Hirschberg. 2025. Learning more with less: Self-supervised approaches for low-resource speech emotion recognition. *arXiv preprint arXiv:2506.02059*. Preprint.

Kristen L. Granger, Melissa Washington-Nortey, Jason C. Chow, Michael D. Broda, Alexandra Montesion, Kevin S. Sutherland, and Maureen A. Conroy. 2024. Child gender and challenging behavior influences early childhood teachers' use of behavior specific praise. *Psychology in the Schools*, 61:4688–4704.

Chanhee Jeong and Aditya Nyamathi. 2026. Causal emotion recognition in conversation: Context saturation and discourse-marker evidence. *Transactions on Machine Learning Research*.

Daniel Kahneman and Dale T. Miller. 1986. Norm theory: Comparing reality to its alternatives. *Psychological Review*, 93(2):136–153.

Bethany U. Kamana, Claudia L. Dozier, Elizabeth A. Foley, and April N. Hubbs. 2022. Descriptive versus general praise with and without edibles for acquisition in young children. *Behavioral Interventions*, 37(2):222–238.

Bidisha Maji, Rishabh Guha, Aurobinda Routray, Shabnam Nasreen, and Debasish Majumdar. 2024. Investigation of layer-wise speech representations in self-supervised learning models: A cross-lingual study in detecting depression. In *Proceedings of Interspeech 2024*. ISCA.

Lucas Martinez-Lucas, Prashanth Mote, Amir R. Naini, Mohammed Abdelwahab, and Carlos Busso. 2026. MSP-Conversation: A corpus for naturalistic, time-continuous emotion recognition. *arXiv preprint arXiv:2603.22536*. Preprint.

Jenna S. Morton, Moïra Mikolajczak, and Olivier Luminet. 2022. New perspectives on the praise literature: Towards a conceptual model of compliment. *Current Psychology*, 41(9):6038–6050.

Claudia M. Mueller and Carol S. Dweck. 1998. Praise for intelligence can undermine children's motivation and performance. *Journal of Personality and Social Psychology*, 75(1):33–52.

Amir R. Naini, Luis Goncalves, A. N. Salman, Prashanth Mote, Ilker R. Ülgen, Thomas Thebaud, Laureano Moro-Velazquez, L. Paola Garcia, Najim Dehak, Berrak Sisman, and Carlos Busso. 2025. The interspeech 2025 challenge on speech emotion recognition in naturalistic conditions. In *Proceedings of Interspeech 2025*. ISCA.

Zahra Omidi and John H. L. Hansen. 2026. Learning from annotation uncertainty: Entropy-aware curriculum for speech emotion recognition. *arXiv preprint arXiv:2606.27536*. Preprint.

Takamura Onishi, Akihiro Yamauchi, Ryo Ishii, Yuki Aono, and Atsushi Miyata. 2020. Analyzing nonverbal behaviors along with praising. In *Proceedings of the 2020 International Conference on Multimodal Interaction*, pages 609–613. ACM.

Rachel Ostrand and Eleanor Chodroff. 2021. It's alignment all the way down, but not all the way up: Speakers align on some features but not others within a dialogue. *Journal of Phonetics*, 88:101074.

Silke Paulmann and Netta Weinstein. 2025. Motivating tones to enhance education: The effects of vocal awareness on teachers' voices. *British Journal of Educational Psychology*, 95(2):551–564.

Amy S. Polick, James E. Carr, and Nicole M. Hanney. 2012. A comparison of general and descriptive praise in teaching intraverbal behavior to children with autism. *Journal of Applied Behavior Analysis*, 45(3):593–599.

Ashley Rila, Allison L. Bruhn, and Lu Wang. 2024. The disaggregated effects of visual performance feedback on teachers' use of behavior specific-praise and reprimands. *Journal of Behavioral Education*, 33:53–80.

Richard M. Ryan and Edward L. Deci. 2000. Intrinsic and extrinsic motivations: Classic definitions and new directions. *Contemporary Educational Psychology*, 25:54–67.

Xiaoran Shi, Ximin Li, and Tomoki Toda. 2025. Speaker-aware multi-task learning for speech emotion recognition. In *Proceedings of Interspeech 2025*. ISCA.

Bart Soenens and Maarten Vansteenkiste. 2020. Understanding the complexity of praise through the lens of self-determination. In Eddie Brummelman, editor, *Psychological Perspectives on Praise*. Routledge.

Colleen Stevens, Tina M. Sidener, Sharon A. Reeve, and David W. Sidener. 2011. Effects of behavior-specific and general praise on acquisition of tacts in children with pervasive developmental disorders. *Research in Autism Spectrum Disorders*, 5(1):666–669.

Nathan Stout. 2020. On the significance of praise. *American Philosophical Quarterly*, 57(3):215–226.

Zachary Tamir, Thomas Thebaud, Jesus Villalba, Najim Dehak, and Oren Kurland. 2025. Multimodal emotion diarization: Frame-wise integration of text and audio representations. In *Proceedings of Interspeech 2025*. ISCA.

Lucas H. Ueda, João G. T. Lima, and Paulo D. P. Costa. 2026. Crab: Multi layer contrastive supervision to improve speech emotion recognition under both acted and natural speech condition. *arXiv preprint arXiv:2603.23673*. Preprint.

Deepak Uniyal and Vinay Abrol. 2025. From pretraining to performance: Benchmarking self-supervised speech models for interspeech-25 SER challenge. In *Proceedings of Interspeech 2025*. ISCA.

Netta Weinstein, Maarten Vansteenkiste, and Silke Paulmann. 2020. Don't you say it that way! experimental evidence that controlling voices elicit defiance. *Journal of Experimental Social Psychology*, 88:103949.

Jessica R. Weyman and Josephine R. Sy. 2018. Effects of neutral and enthusiastic praise on the rate of discrimination acquisition. *Journal of Applied Behavior Analysis*, 51(2):335–344.

Susanne Wolff and Andre Brechmann. 2015. Carrot and stick 2.0: The benefits of natural and motivational prosody in computer-assisted learning. *Computers in Human Behavior*, 43:76–84.

Tianhao Zhang and Zhixing Tan. 2025. ECERC: Evidence-cause attention network for multi-modal emotion recognition in conversation. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. ACL.

Appendix

Table A1: Inter-annotator agreement on the linguistic rubric, before and after the rubric revision round.

Construct	Round 1	Round 2
Praise detection	0.803	0.854
Specificity (any criterion)	0.743	0.797
Verb (C1)	0.652	0.758
Echo (C2)	0.749	0.914
Personal (C3)	0.711	0.841

Table B1: Optimization and evaluation.

Setting	Value
WavLM checkpoint	WavLM-Large, layers 0 to 8 of 24, frozen
Praat	13 hand-crafted prosodic features per praise segment (librosa and parselmouth)
RoBERTa	774-d per turn, 768 from roberta-base plus 6 structural features
Frame encoder	1-layer unidirectional GRU
Turn-level encoder	1-layer GRU, or 1-layer Transformer (4 heads, FF 128) in SERMA-Transformer
Dropout	0.1
Optimizer	AdamW, weight decay 1e-4
Learning rate	1e-3
Batch size	16
Epochs	40, early stopping patience 5
Loss	mean of the two head losses, class-weighted CE with inverse-frequency weights fit on the training fold (BCE for binary tasks)
Seeds	base 42, offset by fold index, one run per fold

Prompt

Praise detection prompt
What is praise?
Praise is a positive evaluative statement — the teacher signals "I approve of what you did", not merely "I heard you" or "that is correct."
Decision — apply in order, stop at the first match
1. Contains a recognized praise expression → PRAISE
Single words: awesome, excellent, fantastic, wonderful, beautiful, brilliant, amazing, outstanding, terrific, bravo, phenomenal, impressive, incredible, marvelous, fabulous, perfect (evaluating work).
Phrases: good/great/nice + job/work/answer/thinking/try/reading/effort, well done, way to go, keep it up, I'm proud of you, I'm impressed, love it, that's + [any single-word above].
2. Explicitly validates what the student did or how they did it → PRAISE
e.g. "You read that correctly", "I can tell you practiced", "That's a clever way to think about it".
3. A short acknowledgment/confirmation/politeness word combined with any positive evaluative word → PRAISE
e.g. "OK, perfect", "That's right, good", "Thank you, nice job".
4. Otherwise → NO PRAISE. In particular, these alone are NOT praise:
- Bare acknowledgments: good, great, nice, right, yes, yeah, ok, all right, sure, mm-hmm, come on.
- Bare confirmations: that's right, correct, exactly, you got it right.
- Politeness / procedural: thank you, please, your turn, turn to page X, let's move on.
- Clearly sarcastic uses of positive words.
Boundary cases
- "Good." alone → NO PRAISE. "Good job." → PRAISE.
- "That's right." alone → NO PRAISE. "That's right, good." → PRAISE (Rule 3).
- "Perfect." after a student's answer → PRAISE. "Perfect, now sit down." → still PRAISE (the evaluative use came first).
- Praise + a following instruction still counts as praise ("Nice job. Turn to page 5." → PRAISE).
- "You read the sentence." → NO PRAISE (neutral). "You read the sentence correctly." → PRAISE (evaluative).

Praise classification prompt
Task
The teacher utterance below already contains praise. Decide which of the three specificity criteria are satisfied. Any one criterion is enough to make the praise SPECIFIC.

Return a list of the satisfied criterion numbers (e.g. [1, 3], or [] if none).
The three criteria
Criterion 1 — Descriptive expression.
The praise clause contains a word or phrase (verb, gerund, noun, or noun phrase) that names what the student did or produced.
✓ "Nice reading." · "Good job listening." · "Great answer." · "Beautiful handwriting." · "That's a creative solution."
× "Good job." · "Awesome." · "Perfect." (generic praise words alone)
Criterion 2 — Echo or paraphrase.
The teacher repeats, restates, or paraphrases a content word/phrase from the student's prior turn. Requires a prior student turn to check.
✓ Student: "gravity" → Teacher: "Gravity — exactly right!"
✓ Student: "I added 5 and 3" → Teacher: "You combined them — perfect!"
× Teacher: "That's great." with no matching content word from the student.
Criterion 3 — Student name attached to the praise.
The student's name is directly connected to the praise clause by comma (vocative). Pronouns like you/your do NOT count.
✓ "Good job, Sarah." · "Sarah, that was excellent!"
× "Good job. Sarah, read this." (name belongs to the following instruction, not the praise)
Key boundary rule
A descriptive word or name only counts if it is inside the praise clause, not inside a following instruction or question.
- "Nice job explaining that." → C1 ✓ (attached)
- "Nice job. Could you explain this?" → C1 × ("explain" is in a new request)
- "Good job, Sarah." → C3 ✓
- "Good job. Sarah, read this." → C3 ×

Does an Automated Source Integration Score Respond to Revision? Extending Evidence of Construct Validity with a Feedback Experiment

Andrew Potter[1], Yu Tian[1], Kaitlin Van Houghton[1], Manmeet Singh[1], Renu Balyan[2], Maria Goldshtein[1], Laura K. Allen[3], Danielle S. McNamara[1]

[1]Learning Engineering Institute, Arizona State University, USA
[2] School of Arts and Sciences, SUNY Old Westbury, USA
[3]Department of Educational Psychology, University of Minnesota, USA

Abstract

This study examined whether an automated source integration measure responds to revision. In a randomized experiment, undergraduates drafted and revised a source-based essay. Source integration scores increased across drafts, although gains did not differ significantly between students who received source integration feedback and those who did not. Findings support use of the measure for formative evaluation and feedback.

1 Introduction

Source integration (i.e., the coordination, attribution, and synthesis of ideas across multiple documents) is central to academic writing and difficult to assess reliably (McCarthy et al., 2022; Vandermeulen et al., 2023). Automated writing evaluation (AWE) tools have largely not addressed connections between reading and writing (Allen et al., 2024). To address this gap, Potter, Tian, et al. (2026c) modeled source integration as a latent construct estimated from linguistic indices. The present study builds on that evidence by asking whether the measure registers what students change when they revise.

According to Kane (2013, 2016), a score is validated by how well the evidence supports the claims made for it. The claim made for this measure is that it can provide formative feedback to help students revise their source-based writing and improve source integration. Such a claim has two inferences: that the score is accurate and that the feedback derived from it is useful when students decide what to revise (Ranalli et al., 2017). The validation evidence described above addresses the first, and grounding the score in interpretable linguistic features also meets the expectation that automated scores be transparent and construct relevant (Ferrara and Qunbar, 2022). The second calls for a different kind of evidence, drawn from students revising with the score, as when draft-to-draft improvement aligned with feedback received supported a formative validity argument for an AWE system (Correnti et al., 2022). Thus, the present study extends the validity argument by examining how the score behaves when undergraduates revise a source-based essay with formative feedback delivered under one of four randomly assigned conditions.

2 Prior Validity Evidence

Potter, Tian, et al. (2026c) developed the source integration measure for formative use in the Writing Analytics Toolkit (WAT). The development corpus comprised 2,643 source-based argumentative essays written by undergraduates and adult learners across eight prompts and scored by trained raters on a four-point source use and inferencing rubric, with acceptable interrater agreement (quadratic weighted $\varkappa = .74$). Three claims were advanced and supported.

Construct representation. Source integration can be represented as a coherent latent construct rather than individual linguistic features. Of 41 candidate linguistic indices extracted with WAT, 11 were retained across four families, namely citation behavior, quotation behavior, verbatim overlap with the source texts, and semantic similarity with the source texts. A hierarchical confirmatory factor analysis fit with lavaan (Rosseel, 2012) modeled citations, quotations, and overlap as first-order factors; word2vec similarity loaded directly onto a second-order source integration factor. Fit was acceptable, $\chi^2(48) = 900.34$, CFI = .945, RMSEA = .082, and the factor predicted expert human ratings of source use and inferencing strongly, $\beta = .81$, $R^2 = .65$.

Generalizability. The construct functions consistently across prompts and populations. Measurement invariance testing across the development corpus and PERSUADE 2.0 (Crossley et al., 2024)

Proceedings of the Artificial Intelligence in Measurement and Education Conference (AIME-Con) – Volume 2: Works in Progress, pages 483–490
October 5-7, 2026 ©2026 National Council on Measurement in Education (NCME)

supported metric but not scalar invariance, indicating that factor loadings held across corpora while intercepts differed. The measure therefore supports the claim that the factor can be applied to novel essays.

Predictive utility. The construct improves prediction of human ratings. The strongest agreement came from concatenating the factor score with BERT embeddings of the essay (Devlin et al., 2019), $\varkappa$ = .43, quadratic weighted $\varkappa$ (QWK) = .68. That combination outperformed embeddings alone ($\varkappa$ = .39, QWK = .65), feature-based machine learning (QWK = .45 to .47), and rubric-driven GPT-5.2 scoring (QWK = .58).

In general, these claims establish that the measure represents its construct, can be generalized across tasks, and reflects judgement of expert trained human raters. All three bear on the accuracy inference. The second inference, about usefulness when students revise, is addressed in the present study.

3 The Present Study

The present study extends prior validation work by examining whether automated source integration scores support an inference required for formative evaluation and feedback, that draft-to-draft score change can be interpreted as change in students' source integration. Following Kane (2013, 2016), we treat this as an interpretation and use claim and ask how strongly the available evidence supports it.

Evidence is stronger if scores increase from first to revised draft, if that increase remains after accounting for changes in essay length and overall writing quality, and if gains are larger among students whose feedback included source integration. The covariate adjustment addresses two rival interpretations, that score change reflects essay expansion or general writing improvement. Feedback condition corroborates rather than isolates, because source integration appeared alongside three other metrics.

RQ1. Responsiveness to revision and feedback context. To what extent do source integration scores increase from first to revised draft, does that increase remain after accounting for changes in essay length and overall writing quality, and does it differ between students whose feedback included source integration and students whose feedback did not?

RQ2. Comparison of score versions. Prior work compared the factor score and the hybrid score on single drafts, so this question asks whether the evidence for the claim depends on which is used. Do the two versions differ in responsiveness to revision, sensitivity to feedback context, and resistance to changes in essay length and overall writing quality?

4 Method

4.1 Participants

Participants were 138 undergraduates at a large public university in the southwestern United States who took part for course credit (M age = 20.04, SD = 1.56, n = 137 reporting; 57.2% men, 42.8% women; 50.0% White, 21.7% multiracial, 13.0% Hispanic or Latino, 10.9% Asian, 4.4% Black or other; 33.3% grew up speaking a language other than or in addition to English). They were recruited through the university research participation system and randomly assigned to one of four between-subjects feedback conditions. Chi-square tests and one-way analyses of variance confirmed no condition differences in demographic or self-efficacy measures. Two participants lacked complete data for both drafts, leaving 136 participants in the analytic sample.

4.2 Research Design

The study used a randomized between-subjects design with four feedback conditions. All participants wrote and revised in the same platform and completed identical tasks; conditions differed only in the feedback available during revision. Participants were assigned to a control condition (n = 36), in which automated metrics were disabled; a WAT-only condition (n = 33), in which students saw platform metrics; a chatbot-only condition (n = 37), in which students received conversational feedback from a large language model; or a combined condition (n = 32), in which students received both forms of feedback. After excluding two participants with incomplete draft data, the analytic sample included 35 control, 33 WAT-only, 36 chatbot-only, and 32 combined-condition participants. Chi-square tests and one-way analyses of variance indicated no condition differences in demographics.

4.3 Materials

Writing task. Students wrote a source-based argumentative essay on teacher retention, taking a

position on why teachers leave the profession and how the problem should be addressed and supporting that position with two sources adapted from online publications, a think-tank blog on national teacher shortages and a newspaper article on teacher burnout. The topic is not represented among the eight prompts of the corpus used to develop and validate the measure, so it is a genuinely unseen task for the scoring model.

WAT-Classroom. WAT-C (McNamara et al., 2026) is a web-based writing platform that provides automated analytics on student writing. Four metrics were enabled for this study, namely source integration, language variety, development of ideas, and academic writing, each shown as a percentile with a short description of what the metric captures and what a writer might do to strengthen it. The panel also displayed a holistic performance score and a word count. The source integration percentile is the rank transform of the same latent factor score evaluated here, so students in the WAT-feedback conditions saw, and could act on, one of the two score versions under comparison. The remaining three metrics are derived by principal component analysis of 72 linguistic features computed over a corpus of source-based essays scored by trained expert human raters (Potter, Crossley, et al., 2026). See Potter, Serhan, et al. (2026) for the design of the panel and the wording of its feedback.

WAT-Chatbot. A pilot conversational agent built with ASU CreateAI Builder on GPT-5.2 (temperature = 0.5) complemented the platform. Its system prompt directed it to offer revision-oriented strategies, to emphasize accurate source use, clarity, and argumentation, and to redirect any request to rewrite text back to the student's own revision, scaffolding revision without off-loading the writing.

4.4 Procedure

Participants completed the study asynchronously in a single online session of approximately 85 minutes. After an orientation to the feedback tools available in their condition, they wrote a first draft with no access to any tool, practiced by revising a peer essay using their condition's feedback, then revised their own essay and completed a survey. Because the first draft was written without feedback and the second was written with it, draft-to-draft change indexes revision under the assigned condition.

4.5 Measures

Source integration. All 272 essays (136 students × 2 drafts) were processed through WAT and scored twice. The factor score was obtained by applying the published scaler and the confirmatory model described in Section 2. The hybrid score was obtained from the architecture that performed best in the earlier predictive comparison, a BERT model that concatenates essay embeddings with the four factor scores, trained on all eight prompts of the 2,643-essay development corpus and producing a value on the scale of the averaged human rating (Potter, Tian, et al., 2026).

Because teacher retention is a new topic, no human source integration ratings exist for these essays and scoring accuracy on this prompt cannot be observed directly. We estimated it by simulating the same situation inside the development corpus, training models on six prompts and evaluating them on the two prompts they had never seen (see Potter, Tian, et al., 2026, for details). Agreement with human ratings was QWK = .64 on one held-out prompt and .43 on the other (pooled QWK = .61, r = .72, n = 791), below the QWK = .68 reported for a random split in which every prompt appears during training. These results suggest that cross-prompt transfer is more difficult than random-split estimates imply, so the pooled held-out agreement provides the more conservative benchmark for the present study. The held-out models were used only to estimate expected transfer accuracy, and all study essays were scored with the final model trained on the full eight-prompt development corpus.

The two score versions are computed on different scales; therefore, their raw values are not directly comparable. The hybrid score was trained on the averaged human rating, which ranged from 0.5 to 4.0 in half-point steps across the 2,643 development essays. That target was min-max normalized for training, and predictions were mapped back to the original scale; thus, hybrid values are reported in rubric points. In this sample, the hybrid score ranged from 1.14 to 3.33 ($M = 2.64$, $SD = 0.35$), and a change of 0.10 rubric points is roughly one-third of a within-sample standard deviation. The factor score is a regression factor score estimated from indices standardized against the development corpus. It has no rubric anchor and, like all regression factor scores, is shrunken toward the corpus mean, which is why its dispersion is narrow (range = -0.56 to 1.42, $M = 0.06$, $SD = 0.20$). Raw

values were analyzed for both versions, and standardized change (d_z) is the metric on which the two can be compared directly.

Covariates. Both drafts were also scored by WAT-C for holistic writing quality and word count. WAT-C holistic scores show strong agreement with expert-trained human raters on source-based writing tasks (McNamara et al., 2026).

Revision behavior. Edit operations between drafts were computed with TRACE (Tian et al., in preparation), which estimates text change at the span, sentence, and token levels and returns counts of insertions, deletions, replacements, and moves. Revision counts were available for 133 participants and were right-skewed (total operations M = 15.31, SD = 20.25, Mdn = 8, range 0 to 134).

4.6 Data Analysis

All analyses were conducted separately for the factor score and the hybrid score. Draft-to-draft change was tested with paired-samples t tests and, because change scores were not normally distributed, with Wilcoxon signed-rank tests on the same pairs. To examine whether score change reflected source integration rather than general essay expansion or improvement, we regressed score change on change in essay length and change in holistic writing quality. The adjusted intercept estimated the mean gain remaining after both covariates were accounted for. Feedback context was tested with an analysis of covariance predicting the revised-draft score from the first-draft score and whether the student received feedback that included source integration.

5 Results

5.1 RQ1. Responsiveness to Revision and Feedback Context

Scores increased from first to revised draft on both versions. On the hybrid score, essays gained $M = 0.076$ rubric points, 95% CI [0.045, 0.107], $t(135) = 4.82$, $p < .001$, $d_z = 0.41$, with the rank test in agreement, $z = -5.23$, $p < .001$. On the factor score, the two tests diverged: the paired-samples t test was not significant, $M = 0.017$ factor-score units, $t(135) = 1.46$, $p = .148$, whereas the Wilcoxon test was significant, $z = -4.25$, $p < .001$.

The divergence between the parametric and rank-based tests was traced to one influential participant. The participant's first draft was 538 words, 469 of which (87%) appeared in properly attributed quotations, indicating extensive but credited source use rather than undisclosed copying. The revision reduced the essay to 277 words and removed every quotation; the factor score decreased by 1.34 factor-score units, relative to a sample standard deviation of 0.13. That value inflated the standard deviation on which the parametric test depends, whereas the rank-based test was less sensitive to it. The exclusion was not prespecified, and because the case is substantively interpretable rather than an artifact, Table 2 reports each test with and without it. Excluding this participant, both score versions increased under both tests, hybrid $t(134) = 5.42$ and factor $t(134) = 4.70$, both $p < .001$. Results below are reported for the full sample, and the exclusion is noted only where it changes a conclusion.

Score gains were partly, but not wholly, associated with essay expansion and overall writing quality. Revised drafts were longer (M = +18.4 words) and received higher holistic scores (M = +0.06). Score change correlated moderately with both length and quality change for the hybrid score (r = .42 and .40) and somewhat less strongly for the factor score (r = .31 and .27), indicating that part of the gain was shared with general essay improvement. After simultaneous adjustment for change in length and holistic quality, the hybrid gain remained significant, +0.058, t(131) = 3.91, p < .001, roughly three-quarters of the unadjusted effect. The factor gain did not remain significant in the full sample, +0.007, t(131) = 0.62, p = .539, although it did after the influential case was excluded, +0.024, t(130) = 4.06, p < .001.

There were no significant differences in source integration scores by feedback condition. Table 1 reports change by condition. Controlling for the draft-1 score, students whose feedback included source integration did not gain more than students whose feedback did not on the hybrid score, $b = 0.033$, $t(133) = 1.16$, $p = .249$. On the factor score, the coefficient was marginal, $b = 0.035$, $t(133) = 1.84$, $p = .067$, and reached significance only when the influential case was excluded. Score change was also unrelated to how much students revised. Edit operations between drafts, computed with TRACE (Tian et al., in preparation), were right-skewed ($M = 15.31$, $SD = 20.25$, $Mdn = 8$, range = 0 to 134, $n = 133$) and did not predict score change on either version (hybrid $r = -.02$, factor $\rho = .15$).

The feedback result should be interpreted cau-

tiously for three reasons. First, source integration was one of four metrics on the panel, so students were not directed specifically toward source use and any construct-specific effect may have been diluted. Second, with group sizes of 65 and 71, the design had 80% power to detect $d = .48$, whereas the observed adjusted effects were smaller ($d = .20$ and $d = .32$). The study was therefore underpowered for effects of the observed size. Third, the largest gain occurred in the combined condition ($+0.128$), but the control condition also gained substantially ($+0.090$), a pattern more consistent with revision itself driving change than with any particular feedback source.

The core evidence supports the claim that score change can be interpreted as revision-related change in source integration. Scores increased when students revised, and, for the primary measure, the increase remained after adjustment for essay length and overall writing quality. The feedback comparison did not provide corroborating evidence. Thus, the evidence supports reading score change as change in source integration, but not attributing that change to the feedback students received.

5.2 RQ2. Comparison of Score Versions

The two score versions were related but not redundant. They correlated moderately on draft-1 scores (r = .63), draft-2 scores (r = .61), and change scores (r = .59), leaving substantial independent variance.

For responsiveness to revision, the hybrid score detected change under both the parametric and rank-based tests in the full sample, whereas the factor score did so only under the rank-based test. For resistance to construct-irrelevant change, the hybrid gain remained significant after adjustment for essay length and holistic quality, whereas the factor gain did not. For sensitivity to feedback context, neither score version showed an effect that held across specifications.

The difference between the score versions was driven by the factor score's sensitivity to one influential case. Removing one participant in 136 halved the factor score's change-score standard deviation from 0.135 to 0.066, increased its paired t from 1.46 to 4.70, changed its adjusted gain from p = .539 to p < .001, and changed its condition test from p = .067 to p = .005. In contrast, the hybrid score shifted only from t = 4.82 to 5.42, and its condition test remained nonsignificant.

The case illustrates why the factor score was less stable. The factor score is computed from citation, quotation, and overlap indices. A draft composed of 87% quoted text can therefore score near the top of the distribution on quotation and overlap, even if that material is not well integrated into an argument. This limitation was anticipated in the validation study, where overlap indices loaded positively on the construct despite their theoretical association with weaker integration, because such indices cannot distinguish unproductive copying from legitimate quotation (Potter, Tian, et al., 2026). The hybrid score also incorporates the essay text, and its estimate for this student fell by 0.63 rubric points, compared with a 1.34-unit drop on the factor score.

One asymmetry should be interpreted cautiously. Students saw the factor score on the feedback panel, and only the factor score's condition coefficient approached, then crossed, the conventional significance threshold after the influential case was removed. Because the effect was small and depended on one case, we treat it as a direction for further study rather than as evidence of a feedback effect.

Overall, both score versions provided evidence of responsiveness to revision, but the hybrid score provided more stable evidence across the full set of tests. The evidence for revision-related change did not depend strongly on score version once the influential case was set aside. However, the evidence for resistance to construct-irrelevant change and sensitivity to feedback context did depend on score version, with the factor score's conclusions turning on a single participant.

6 Discussion

Potter, Tian, et al. (2026) established evidence of construct validity for the source integration measure by demonstrating that it represents its construct, functions consistently across corpora, and predicts expert human ratings. However, that evidence concerned single drafts. The present study extends this validity evidence to a second occasion of measurement in which students revised a source-based essay with or without feedback on source integration. Scores increased when students revised, and, for the hybrid score, the increase remained after adjustment for changes in essay length and holistic writing quality. These findings support interpreting draft-to-draft score change as change in students' source integration, an inference required for use in formative evaluation and feedback.

However, gains in source integration scores were not larger among students whose feedback included source integration. Three design features qualify this finding. First, source integration appeared alongside three other metrics, so students were not directed specifically toward source use. Second, the study was underpowered for effects of the observed size. Third, scores increased in every condition, including the control condition, suggesting that revision itself accounted for much of the change. Isolating a construct-specific feedback effect would require a design in which source integration feedback is the only feedback feature that varies. Future experimental studies should isolate source integration feedback from other feedback dimensions, vary its presence systematically, and test whether students who receive targeted source integration feedback show larger score gains than students who receive either no automated feedback or feedback on other writing dimensions.

Furthermore, the method in which the source integration score was calculated diverged in ways that are important for future use in AWE systems. One participant in 136 moved the factor score's adjusted gain across the conventional significance threshold, while leaving the hybrid score largely unchanged. That participant quoted sources extensively in the first draft and removed all quotations in the revision. This pattern is consistent with Potter, Tian, et al. (2026), who noted that overlap indices loaded positively on the construct because they cannot distinguish unproductive copying from legitimate quotation. Indices counted without reference to the surrounding text can register a heavily quoted draft as strong source use, whereas the essay text contains information those counts omit. These results extend the earlier finding that the hybrid architecture best predicted human ratings by showing that its advantage also appears across occasions of measurement.

6.1 Limitations

Several limitations qualify the findings and point to future research. Human source integration ratings were unavailable for the study essays, so scoring accuracy on this prompt was estimated through cross-prompt transfer rather than observed directly. In addition, the covariates used to evaluate rival interpretations, essay length and holistic writing quality, were automated. These constraints limit conclusions about scoring accuracy on the present prompt.

The design also limits generalizability. Students revised once, on one prompt, therefore the findings are related to a single revision occasion rather than to development over time. Future studies should follow writers across multiple source-based tasks and revision cycles to test whether score change accumulates in the direction instruction intends.

7 Conclusion

The validity argument for this measure now includes evidence from both single-draft scoring and revision-based use. Potter, Tian, et al. (2026) established that source integration can be represented as a coherent latent construct, that the construct generalizes across prompts and populations, and that it improves prediction of human ratings, especially when combined with essay embeddings. The present study adds evidence that draft-to-draft change in the score can be interpreted as change in students' source integration. This interpretation is qualified by the absence of corroborating evidence from feedback condition and by the fact that the evidence was stronger and more stable for the hybrid score than for the factor score.

Two implications follow for practice. First, responsiveness to revision is a distinct source of validity evidence for scores intended to support formative evaluation and feedback, and it should be demonstrated rather than inferred from single-draft accuracy. Second, decisions about score deployment should consider how candidate score versions behave across drafts, because formative interpretation depends on score change during revision.

Acknowledgments

The research reported here was supported by the Institute of Education Sciences, U.S. Department of Education, through Grant R305A180261 to Arizona State University. The opinions expressed are those of the authors and do not represent the views of the Institute or the U.S. Department of Education.

References

Laura K. Allen, P. Öncel, and L. E. Flynn. 2024. Expanding AWE to incorporate reading and writing evaluation. In M. D. Shermis and J. Wilson, editors, *The Routledge International Handbook of Automated Essay Evaluation*, pages 575–589. Routledge.

R. Correnti, L. C. Matsumura, E. L. Wang, D. Litman, and H. Zhang. 2022. Building a validity argument for

an automated writing evaluation system (eRevise) as a formative assessment. *Computers and Education Open*, 3:100084.

S. A. Crossley, P. Baffour, Y. Tian, A. Franklin, M. Benner, and U. Boser. 2024. A large-scale corpus for assessing written argumentation: PERSUADE 2.0. *Assessing Writing*, 61:100865.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics*, pages 4171–4186. Association for Computational Linguistics.

S. Ferrara and S. Qunbar. 2022. Validity arguments for AI-based automated scores: Essay scoring as an illustration. *Journal of Educational Measurement*, 59(3):288–313.

M. T. Kane. 2013. Validating the interpretations and uses of test scores. *Journal of Educational Measurement*, 50(1):1–73.

M. T. Kane. 2016. Explicating validity. *Assessment in Education: Principles, Policy & Practice*, 23(2):198–211.

K. S. McCarthy, E. F. Yan, L. K. Allen, A. N. Sonia, J. P. Magliano, and D. S. McNamara. 2022. On the basis of source: Impacts of individual differences on multiple-document integrated reading and writing tasks. *Learning and Instruction*, 79:101599.

D. S. McNamara, A. Potter, Z. Serhan, M. Singh, N. Patne, T. Arner, R. D. Roscoe, L. K. Allen, and S. A. Crossley. 2026. The writing analytics tool: A learning engineering approach to designing AI-supported writing instruction. In *Proceedings of the Learning Engineering Research Network Convening (LERN 2026)*. EdTech Archives.

A. H. Potter, S. A. Crossley, P. Öncel, L. A. Allen, X. Xu, R. Balyan, T. Arner, R. D. Roscoe, and D. S. McNamara. 2026a. From words to components: Mapping linguistic features in source-based writing. *Assessing Writing*, 70:101108.

A. H. Potter, Z. Serhan, N. A. Patne, P. Öncel, I. Ahmed, T. Arner, R. Islam, R. D. Roscoe, L. K. Allen, S. A. Crossley, and D. S. McNamara. 2026b. Human-AI collaboration for qualitative analysis in participatory design: Refining the writing analytics tool. *Journal of Educational Data Mining*, 18:113–155.

A. H. Potter, Y. Tian, R. Balyan, M. Goldshtein, L. K. Allen, and D. S. McNamara. 2026c. Source integration in source-based writing: From theoretical modeling to predictive analytics. *Assessing Writing*, 69:101081.

J. Ranalli, S. Link, and E. Chukharev-Hudilainen. 2017. Automated writing evaluation for formative assessment of second language writing: Investigating the accuracy and usefulness of feedback as part of argument-based validation. *Educational Psychology*, 37:8–25.

Y. Rosseel. 2012. lavaan: An R package for structural equation modeling. *Journal of Statistical Software*, 48(2):1–36.

Y. Tian, K. Christhilf, A. Potter, S. Graham, J. Early, and D. S. McNamara. in preparation. TRACE: Automated multi-granularity analysis of text revision. Manuscript in preparation.

N. Vandermeulen, E. Van Steendam, and G. Rijlaarsdam. 2023. Introduction to the special issue on synthesis tasks: Where reading and writing meet. *Reading and Writing*, 36(4):747–768.

Table 1

Mean draft-1 to draft-2 change by feedback condition. Ns are the analyzed sample.

Condition	n	Hybrid Δ	Factor Δ
Control	35	$+0.090$	$+0.025$
WAT-Only	33	$+0.053$	-0.004
Chatbot-Only	36	$+0.037$	$+0.006$
WAT+Chatbot	32	$+0.128$	$+0.041$
Overall	136	$+0.076$	$+0.017$

Table 2

Change, adjustment, and condition tests for both score versions, on the full sample and excluding the influential case.

	Hybrid		Factor	
Statistic	Full	$-$Out	Full	$-$Out
Change M	$+.076$	$+.081$	$+.017$	$+.027$
Change SD	.184	.174	.135	.066
Paired t	4.82^{***}	5.42^{***}	1.46	4.70^{***}
Wilcoxon z	-5.23^{***}	-5.44^{***}	-4.25^{***}	-4.45^{***}
d_z	.41	.47	.12	.40
Adjusted b	$+.058$	$+.064$	$+.007$	$+.024$
Adjusted t	3.91^{***}	4.41^{***}	0.62	4.06^{***}
ANCOVA b	$+.033$	$+.042$	$+.035$	$+.032$
ANCOVA t	1.16	1.53	1.84	2.88^{**}

Note. $-$Out = excluding the influential case. Hybrid values are in rubric points and factor values are in factor-score units, so only d_z is comparable across versions. Adjusted rows give the intercept from regressing score change on change in essay length and change in holistic quality. ANCOVA coefficients are for source integration feedback controlling the draft-1 score. $^{**}p < .01$. $^{***}p < .001$.

Author Index

Abulela, Mohammad A. A., 184
Afrin, Tazin, 100
Agnoli, Sam, 211
Ahmadli, Shahana, 397
Allen, Laura K., 483
Anifowose, Charles, 315
Antal, Judit, 140
Arroyo, Ivon, 289

Babalola, Damilola, 358
Bailey, Paul D., 62
Balyan, Renu, 483
Barber, Justin, 405
Barrett, Michelle, 1
Bediwy, Ahmed H., 193, 200, 382
Bellows, Martha, 159, 193, 200, 322
Buddarapu, Naga, 358
Buzo-Casanova, Enrique, 37, 86

Chan, Cyrus Sai-Cheong, 315
Chandler, Chelsea, 350
Chandrashekar, Tharunya, 130
Chauhan, Kargi, 130
Chen, Jianshen, 140
Cheng, Britte Haugan, 329
Choi, Jinah, 1
Choi, Jinnie, 145
Christhilf, Katerina, 45
Christie, Thomas, 168
Chrzan, Michael Leon, 24
Codoreanu, Alex, 79
Corrigan, Brittany, 275
Cruz Juárez, Adriana, 37
Cutumisu, Maria, 231

D'Brot, Maria, 405
Dallas, Andrew, 275
de Souza, Talita de Paula Cypriano, 397
Deane, Paul, 358
Demmans Epp, Carrie, 231
Deng, Nina, 260
Diana, Nicholas, 464
DiCerbo, Kristen, 329
Dogan, Enis, 282
Domingue, Benjamin, 24
Dulhunty, Mark, 79

Early, Jessica, 45

Fakiyesi, Vincent, 405
Fancsali, Steve, 389
Flor, Michael, 207
Foltz, Peter W, 350
Frost, Julius, 423

García-Lima, Milton, 86
García-Minjares, Manuel, 37, 86
Gattupalli, Sai, 289
Gianopulos, Garron, 152, 177
Golden, Richard, 235
Goldshtein, Maria, 483
Gorney, Kylie, 140
Graham, Steve, 45
Gregg, Nikole, 275
Gushta, Matthew, 184

Ha, Le An, 100
Han, Suhwa, 248
Harik, Polina, 100
Harrold, Brian, 405
Hauru, Markus, 168
Hellman, Scott, 159
Hemenway, Michael, 405
Hemenway, Michael P., 322
Hong, Seong Eun, 377
Hosseinpourkhoshkbari, Reyhaneh, 235
Huang, Ruiping, 448
Huang, Yue, 267

Isotani, Seiji, 397

Jacobs, Gregory M., 193, 200
Jain, Yash Bipin, 235
Jiang, Yanlin, 275

Karimi-Malekabadi, Farzan, 1
Kave, Colleen Burke, 57
Kim, Hyunah, 315
Kim, Minseok, 344, 397
Kim, Seongeun, 93
Ko, Mon-Lin Monica, 350
Kolle, Harvey Ngoe, 231
Kruse, Michael B., 62

Lee, Chansoon, 140
Lee, William, 289

Li, Lanrong, 184
Li, Zhen, 369
Liaqat, Amna, 231
Lim, Sangdon, 152, 177
Lin, Qiao, 57
Liu, Jing, 472
Lottridge, Sue, 159
Lottridge, Susan, 405
Lovelace, Temple, 459
Lu, Yikai, 113
Lynch, Collin, 358
Lyness, Jeffrey Marc, 57
Lyu, Meng, 441

Ma, Wanjing Anya, 207
Maeda, Hotaka, 113, 130
Malkin, Linda, 17
Marigo, Alessia, 17
McGrane, Joshua A, 306
McNamara, Danielle, 45
McNamara, Danielle S., 483
Mee, Janet, 100
Mehta, Shruti, 397
Meng, Lionel, 335
Middlebrook, Kamal, 62
Mitros, Piotr, 358
Mojoyinola, Mubarak O., 382
Moore, Steven James, 464
Moses, Tim, 93
Murphy, April, 389

Nam, Sungjin, 52
Nfonsang, Neba, 459
Norberg, Kole, 389

O'Brien, Kristin, 57
O'Connor, Brendan, 289
O'Reilly, Tenaha, 207
Ogut, Burhan, 282
Olivera-Aguilar, Margarita, 62
Omopekunola, Moses Oluoke, 441
Ortega-Sánchez, Elibidú, 37, 86

Patz, Richard, 369
Perez, Alexandra, 177
Perez, Alexandra Lane, 152
Potter, Andrew, 45, 483
Powers, Sonya, 1

Radwan, Nizam, 315
Rafferty, Anna, 168
Raj, Aleena K, 260
Rao, V. N. Vimal, 216
Rezayi, Saed, 100
Richburg, Aquia, 275, 298
Rijmen, Frank, 248
Rikoon, Samuel H., 62
Ritter, Steve, 389

Schneider, Christina, 152
Schneider, M. Christina, 177
Schneider, Wei Shuang, 69
Scott, Caleb, 358
Seo, Daeryong, 104
Sharif, Melanie, 159
Shen, Qian, 9
Shi, Qingzhou, 211
Shin, David, 57
Singh, Manmeet, 483
Sun, Huaping, 57
Sánchez-Mendiola, Melchor, 37, 86

Tang, Steven, 369
Tian, Yu, 45, 483
Tong, Lingbo, 335
Torres Irribarra, David, 306
Trout, Nick, 93
Tse, Winnie Wing-Yee, 315

Uttal, David, 211

Van Houghton, Kaitlin, 483
Vuong, Mai Que, 397

Walker, Leah, 145
Walker, Marcus, 275, 298
Wang, Feng, 441
Wang, Zuowei, 207
Whitmer, John, 329
Wright, Laura, 17

Xie, Xin, 17
Xu, Paiheng, 472
Xu, Tianyu, 416
Xu, Xinyu, 416
Xue, Kang, 152, 177

Yin, Yue, 448
Yoo, Jiseung, 472

Zhang, Icy, 335
Zhang, Su, 315

Zhang, Susu, 216
Zhou, Michelle, 397
Zoanetti, Nathan, 79
Zu, Jiyun, 207