

Autoscoring Anticlimax: A Meta-analytic Understanding of AI's Short-answer Shortcomings and Wording Weaknesses

Michael Hardy
Stanford University
hardym@stanford.edu

Abstract

Automated short-answer scoring lags other LLM applications. We meta-analyze 890 culminating results across a systematic review of LLM short-answer scoring studies, modeling the traditional effect size of Quadratic Weighted Kappa (QWK) with mixed effects metaregression. We quantitatively illustrate that the level of difficulty for human experts to perform the task of scoring written work of children has no observed statistical effect on LLM performance. Particularly, we show that some scoring tasks measured as the easiest by human scorers were the hardest for LLMs. Whether by poor implementation by thoughtful researchers or patterns traceable to autoregressive training, on average decoder-only architectures underperform encoders by 0.37—a substantial difference in agreement with humans. Additionally, we measure the contributions of various aspects of LLM technology on successful scoring such as tokenizer vocabulary size, which exhibits diminishing returns—potentially due to undertrained tokens. Findings argue for systems design which better anticipates known statistical shortcomings of autoregressive models. Finally, we provide additional experiments to illustrate wording and tokenization sensitivity and bias elicitation in high-stakes education contexts, where LLMs demonstrate racial discrimination. Code and data for this study are available online.¹

1 Introduction and motivation

Automated scoring of short-answer responses remains stubbornly hard (Ormerod, 2022; Ormerod and Kwako, 2024; Chamieh et al., 2024; Kim et al., 2026; Jurenka et al., 2024). Unlike tasks where large language models (LLMs) have leapt forward, autoscoring algorithms must align with rubric-grounded, meaning-dependent determina-

tions, with high reliability and fairness. We investigate short-answer scoring, and demonstrate that these performance lapses in performance are associated with known fundamental challenges for many LLMs (McCoy et al., 2023). Put another way, we investigate the irony of the greatest technological innovation of the last 20 years—Large Language Models—struggling to meet standards that set by some of the earliest language models, before the turn of this century, used in any evaluation capacity were for automated scoring of written work from students (Shermis and Burstein, 2003).




Figure 1: LM Benchmarks over time from (Kiela et al., 2023). Short Answer Scoring of K12 student writing on the ASAP-SAS dataset, in purple, are added using the same scaling: human score is set to 0 and baseline performance to -1, and results, X , are scaled as $(X - \text{Human}) / |\text{Baseline} - \text{Human}|$. Human and 2012 Baseline results are from Shermis. Point (A) represents the best LLM model, a transformer ensemble, (Ormerod, 2022). Point (B) Best GPT model ensemble, a fine-tuned ensemble, (Ormerod and Kwako, 2024). Point (C) [not plotted] is the best implementation using GPT-4 (Jiang and Bosch, 2024) with prompt engineering and would have a value of -1.52 on this scale. While the plot ends in 2025, to the best of our knowledge this chart still reflects SOTA for this task.

The purpose of this study is to support developers and researchers in creating solutions to the most important problem in education technology: measuring meaningful student learning. A cen-

¹https://github.com/hardy-education/SAS_LLM_meta_analysis

tral question is whether current LLMs, especially those trained to autoregress Internet text, possess the right inductive biases to evaluate whether a child is demonstrating learning through text.

Current autoscoring systems often struggle, despite the hope that the powerful generative capabilities of modern LLMs would provide a breakthrough as has been seen on many other language tasks (Fig. 1). This frustrating plateau suggests that simply applying more powerful general-purpose models or tweaking prompts may be an insufficient strategy. A more fundamental inquiry is required.

We hypothesize that the limitations of LLMs in autoscoring are intrinsically linked to their primary training objective: the autoregressive prediction of internet text. This objective optimizes for textual fluency and pattern matching, not for the deep semantic comprehension and inferential reasoning required to evaluate a student’s understanding. We formulate a set of precise, testable hypotheses:

1. Do LLMs exhibit lower performance on tasks that are more meaning-dependent (e.g., literary comprehension and analysis) compared to those that are more fact-based (e.g., science)?
2. Do autoregressive decoder-only architectures (e.g., GPT) underperform relative to encoder-only architectures?
3. Does the size of a model’s tokenizer vocabulary have a predictable, non-linear relationship with its scoring performance?

2 Background

LLMs are language models and do not directly construct or manipulate ideas and thoughts (Mahowald et al., 2023). Rather, they manipulate language tokens, allowing them to communicate ideas in remarkably sophisticated ways. Whether these models can truly understand the meaning of ideas or are only responding successfully from statistical properties of language from their training remains an open question. Regardless, their training does not produce basic reasoning about ideas in GPT models (Kosoy et al., 2023; Yiu et al., 2023), making meaning-based tasks, such as entity tracking in long sequences (Kim and Schuster, 2023; Liu et al., 2023) and negation tracking (Ettinger, 2020), more challenging for GPT models. Additionally, as will be explored further, subtle changes in word choice can overshadow prompt meaning and lead to performance brittleness (Ceron et al., 2024; Tjuatja et al., 2024; Webson and Pavlick,

2022). Tasks where reading for understanding, analyzing language used to convey ideas, and/or manipulating ideas from text are often more meaning-dependence tasks. These are common in K12. As an example, the GPT-4 technical report showed impressive gains on exams designed for humans over GPT-3.5 (OpenAI et al., 2023), including K12 exams such as Advanced Placement (AP) tests. It more than qualified for college credit on all AP tests except on two AP English tests: language and literature, which require examinees to read, analyze, and respond to ideas conveyed within text. In fact, GPT-4’s large scaling over GPT-3.5 did not even improve performance on these tasks, despite the irony of being language models and posting large gains on other AP tests.

2.1 Sensitivity to Tokenization and Wording

Known weaknesses of LLM architectures include sensitivities to tokenization and formatting (Liang et al., 2023; Sclar et al., 2024; Wang et al., 2024b; Xia et al., 2024). Typos and irregular text have been shown to elicit undesirable behaviors (Greshake et al., 2023; Schulhoff et al., 2023; Sclar et al., 2024), as processes used in creating efficient subword tokenization carry with them the biases arising from however that efficiency is achieved (Land and Bartolo, 2024; Wang et al., 2024b). As an additional measure to illustrate token sensitivity, we include another experiment in Appendix A, where we demonstrate how very different documents are generated from only adding 0-2 leading or trailing spaces in the prompt, even with prompt settings unchanged. This method allows us to generate varying responses without generating any differences in meaning. This additional experiment serves to illustrate that the issues discussed and uncovered in this study are not confined to just automatic scoring.

Tokenizer algorithms commonly used include WordPiece (Devlin et al., 2019), Unigram (Kudo, 2018), and Byte Pair Encoding (BPE) (Gage, 1994; Sennrich et al., 2015), all of which can be implemented by a library, such as SentencePiece (Kudo and Richardson, 2018) or tiktoken (OpenAI, 2024), across a training corpus and then made functional by a token vocabulary containing numeric representations for each token recognized. These components serve as the interface between language and the model components of LLMs, i.e., architecture, data, and training task. Related to tokenization, we would expect that GPT models, which are typically

trained on language and not ideas, would be sensitive to variations in language representing the same meaning (Dentella et al., 2023; Liang et al., 2023). Indeed, GPT models have shown high sensitivity to wording in prompting techniques dissimilarly to humans (Tjuaatja et al., 2024) with only subtle changes in text leading to performance differences of up to 76% accuracy points (McCoy et al., 2023; Sclar et al., 2024). This is further demonstrated by how models respond to the inclusion of irrelevant content can drastically change the response (Branch et al., 2022; Kassner and Schütze, 2020; Shi et al., 2023). Yet on tasks where GPT models should be more sensitive to particular words or meanings in their instructions, they can ignore important content (Min et al., 2022; Webson and Pavlick, 2022) and small textual differences that change meaning (Bisbee et al., 2024; Ettinger, 2020). This reemphasizes that they rely on unknown statistical features of language rather than on human-like meaning. Put another way, we train these models to act like us, but not necessarily to think like us.

2.2 Scoring of Student Work

Accurate automated grading of student work is a highly desired application of language models for educators, however, it has not seen the same benchmark saturation as other model evaluations from the current LLM/AI revolution (Fig. 1). State-of-the-art LLM autoscoring models on publicly available datasets are unsatisfying in that they are not transferrable to other applications as they consist of computationally heavy ensembles of several models, each of which are fine-tuned to a single question to be graded (Ormerod and Kwako, 2024; Ormerod, 2022). Without fine-tuning, GPT-models' best performance, only using prompt engineering, including few-shot learning, with GPT-4 or GPT-4.1, was far worse than the baselines of 2012 (Shermis, 2015; Ormerod and Kwako, 2024; Xiao et al., 2024; Kim et al., 2026). In fact, when researchers provided increasing k-shot examples with labels, no GPT models improved and, in many cases, performed worse (Chamieh et al., 2024). Analyzing this performance as a "challenge" dataset suggests that the LLM is part of the problem (Liu et al., 2019). Despite this, educators have already been using these models for scoring student work (Open Innovation Team and Department for Education, 2024). The task of K12 scoring is discerning whether a student's written response demonstrates proficiency of the knowledge and skills assessed.

Thus, small differences in wording or even punctuation can lead to different scores, however, without an understanding of actual changes in meaning in a student response, an LLM may not be sensitive to the critical features of the writing.

Issues in tokenization, format, and wording become more pronounced as the input text deviates more from the training distribution, a deviation that becomes more likely as the age of the writer decreases. Children learning to write often employ phonetic spelling and imperfect word approximations, whose intended meaning may be easily decipherable for a human, but out-of-vocabulary for a GPT model.² We would expect that LMs that are more robust to formatting issues in tokenization would also be more robust to text irregularities as displayed by children's writing, which could be a factor behind the difference between point (A) and point (B) in Figure 1. This is not to say that automated grading by GPT models is impossible, but that the solution is non-trivial and currently not recommended (Koo et al., 2024). Implications being that any edtech resource using such models and asserting the ability to fairly and accurately assess student learning using students' natural language on an educator-provided or AI generated question would need to produce compelling evidence to support such claims, especially if such measures are ever to be used by educational decision-makers (e.g., reporting scores for students or groups to teachers). Degradation of education would come as a direct result of either misrepresenting student learning. And being able to tell whether something is actually helping students is critical to improving the rigor with which the field moves forward (Jurenka et al., 2024; Klahr, 2013; Liu et al., 2024).

2.3 LLM responses matter: Biases in LLMs

Harmful biases against students are not explicitly measurable in the present study, but challenges in scoring student may also contribute to triggering latent harmful biases within LLMs. They are trained on massive text corpora, effectively encoding historical language patterns—a reflection of the world as it *was*, not necessarily as *it should be*. As such, these GPT "imitation engines" (Bender et al., 2021; Yiu et al., 2023) inherit and propagate existing societal biases (Benjamin, 2019; Shieh et al., 2024;

²An example of writing for a third grader with these traits can be found within the prompt in Appendix D, which was used in the prompt of the feedback in Table 1, an example which was generated directly from OpenAI's chat interface.

Table 1: Racial Bias Autoscoring Example: Score and feedback for identical student work using ChatGPT

“White” student ChatGPT’s assigned score: 1 (out of 2)	“Black” student ChatGPT’s assigned score: 0 (out of 2)
Feedback: 1. Spelling and Grammar: There are several spelling and grammar errors throughout the response. For example, “excited” is misspelled as “exited,” “antiqued” should be “questioned,” “youre” should be “your,” “anoyed” should be “annoyed,” “alein’s” should be “alien’s,” and “finnaly” should be “finally.” 2. Punctuation: There are missing commas in several places, such as after “Yes” and “Oh wait,” and inconsistent punctuation usage , such as missing <i>periods at the end of sentences</i> . 3. Capitalization: Some words are not capitalized when they should be, such as “youre” instead of “your” and “alein’s” instead of “alien’s.” 4. Clarity and Coherence: The response lacks clarity and coherence due to the frequent errors, making it difficult to follow the narrative smoothly. Overall, the response demonstrates some attempt to convey the story but is hindered by numerous language errors that impact readability and understanding .	Feedback: 1. Spelling and Grammar: There are numerous spelling and grammar errors throughout the response. For example, “excited” is misspelled as “exited,” “antiqued” should be “questioned,” “youre” should be “your,” “anoyed” should be “annoyed,” “alein’s” should be “alien’s,” “immediately” is misspelled as “imedately,” and “finnaly” should be “finally.” 2. Punctuation: There are missing commas in several places, such as after “Yes,” “Why,” and “Oh wait.” Additionally, there are incorrect punctuation marks, such as multiple ellipses (...) used instead of periods, and missing periods at the end of sentences. 3. Capitalization: Some words are not capitalized when they should be, such as “youre” instead of “your” and “alein’s” instead of “alien’s.” 4. Clarity and Coherence: The response lacks clarity and coherence due to the numerous errors, making it difficult to follow the narrative smoothly. Overall, the response demonstrates signijicant issues with spelling, grammar, punctuation, and capitalization, which greatly affect readability and comprehension .

Differences in ChatGPT autoscoring and student feedback. Identical prompts with one-word perturbation of “White” to “Black” (full prompt texts in Appendix A) result in different scores, with the “White” student receiving the higher score and different feedback to students, and additional details can be found in Appendix D. **highlighted and italicized** text mark score and textual differences. Text in gray is hallucinatory feedback for “Capitalization” which repeats prior generated text.

Xue et al., 2023). This poses a significant challenge for education, where equity is paramount. LLM-perpetuated biases risk reinforcing existing inequities and creating difficult-to-address data proxies (Gonen and Goldberg, 2019; Warr et al., 2024; Warr, 2024) potentially leading to spurious correlations or masking underlying problems during evaluation (Gururangan et al., 2018; Pang et al., 2023; Poliak et al., 2018).

Current LLMs exhibit biases activated by various factors, including dialectal variations like African-American English (AAE) (Deas et al., 2023; Hofmann et al., 2024b) This mirrors known bias issues within public education itself (Baker-Bell, 2020). For example, Hardy (Hardy, 2025, 2024) demonstrated racial bias in LLM evaluations of classroom instruction quality, even when the models were unaware of the teacher’s race.

While instruction-tuning and prompt engineering can refine LLM outputs, the computationally expensive foundational training on massive datasets remains the primary driver of performance (McCoy et al., 2023). Though prompt injection and jailbreaking can manipulate model behavior (Rossi et al., 2024; Schulhoff et al., 2024, 2023). and

system prompts offer some mitigation (Pawelczyk et al., 2024), it is unlikely prompting can reliably eliminate the biases ingrained from internet text.

Detecting biased performance is challenging for individual users, as LLMs often appear reasonable and trustworthy even when inaccurate (Klingbeil et al., 2024; Wen et al., 2024; Zhou et al., 2024, 2023). Concerningly, both children and adults exhibit increased trust with repeated exposure (Kosoy et al., 2024), a phenomenon demanding deeper investigation given observed behavioral changes in post-secondary learning (Abbas et al., 2024; Nie et al., 2024; Zhai et al., 2024).

Given the pervasiveness of these biases and children’s vulnerability, understanding their extent in K-12 applications is critical. We hypothesize that biases persist even with best practices. To demonstrate this, we conducted an experiment illustrating racial bias in a GPT model providing feedback on a 3rd-grade essay. Identical essays, attributed to either a “White” or “Black” student, received differing scores from ChatGPT. The “White” student received a higher score. These results, detailed in Table 1, underscore the urgency of addressing LLM bias in educational contexts. This example does

not require explanation to recognize that it is problematic. It also mirrors other research (Deshpande et al., 2023; Shieh et al., 2024; Warr et al., 2024).

3 Data

We use mixed-effects modeling for meta-regression to control for study heterogeneity, as well as random heterogeneity across the sample of items, the LLMs employed, and the various machine learning practices employed across the studies. In addition to the random effects, we control for the relative difficulty of scoring the items by using the human QWK (Shermis, 2015) as a covariate. The original dataset used by these studies, the ASAP-SAS³ corpus, comprises diverse prompts and scoring rubrics for short answer scoring. Additional information about the datasets and metadataset can be found in Appendix B.

4 Methods

4.1 QWK outcome

We meta-analyzed results from multiple empirical studies in which large language models (LLMs) were used to score short, child-written constructed responses on a common set of $J = 10$ rubric-scored items. Studies frequently evaluated multiple models and multiple deployment regimes (e.g., prompting vs. fine-tuning). The unit of analysis was a single evaluation result (a model–training configuration evaluated on a particular item within a study), yielding $N = 890$ observed agreement coefficients. Details regarding the systematic review that collected this data and overall statistics for the metadataset can be found in Appendix B and Table 6 (see code and data repository online).

Effect size Each observation is a quadratic weighted kappa (QWK) coefficient, $\kappa_w \in [-1, 1]$, comparing the automated score to human scores with quadratic weights on ordinal disagreements. QWK is a reliability-style measure: it penalizes larger score discrepancies more heavily and corrects for chance agreement. In the case of comparisons between each LLM instance and human raters, this reliability measure also effectively acts as a measure of accuracy that has been used for these kinds of tasks for well over a decade (Shermis, 2015; Shermis and Burstein, 2003). Because correlation-like coefficients exhibit non-constant

sampling variance and nonlinearity near the boundaries, we analyze a Fisher z -transformed effect size $y := \text{atanh}(\kappa_w) = \frac{1}{2} \log\left(\frac{1+\kappa_w}{1-\kappa_w}\right)$, which is approximately normal when κ_w is not extremely close to ± 1 .

4.2 Predictors

Each evaluation result includes the following predictors (coded at the level of the model configuration and/or item):

- **Decoder (GPT-style) vs. encoder** (gpt): an indicator for autoregressive/decoder LLMs versus encoder-based architectures.
- **Meaning dependence** (read): an indicator for items requiring more semantic interpretation (as opposed to more form/keyword-aligned scoring). These items are expected to expose brittleness in tokenization and next-token objectives when facing idiosyncratic child language.
- **Tokenization family** (token): a categorical indicator for BPE, Unigram, or WordPiece. We include a coefficient for each family to represent the reference values of the remaining covariates.
- **Vocabulary size** ($|V|$, vocab) and **quadratic term** ($|V|^2$, $I(\text{vocab}^2)$): to capture non-monotone effects expected when vocabularies are too small (over-fragmentation) or too large (under-trained rare tokens).
- **Model size** (logsize): log parameter count (or an equivalent log-scale size proxy) to represent scaling effects.
- **Human benchmark difficulty** (QWK_{hum} , humqwk): QWK among human raters for the same item, Fisher- z transformed. This serves as a principled covariate for item/rubric difficulty as experienced by expert raters.

4.3 Hierarchical meta-regression

The dataset is multi-way clustered: observations are nested within items, studies, models, and training regimes, and many studies report multiple models and multiple training regimes. Treating all N rows as independent overstates precision. Moreover, substantial heterogeneity is expected due to (i) differences in datasets and annotation processes across studies, (ii) model-specific idiosyncrasies, and (iii) item-specific sensitivity to semantics and potentially child-language and handwriting transcription variation. We therefore estimate multi-level meta-regressions with random effects and, in the most conservative specification, allow *item-*

³<https://www.kaggle.com/c/asap-sas>

specific deviations that vary by model, study, and training regime.

Baseline and increasingly controlled specifications We first fit a sequence of frequentist linear mixed models (LMMs) to assess robustness to different random-effects structures. The general form of our models is:

$$y_{ijkl} = \mathbf{X}_{ijkl}^\top \boldsymbol{\beta} + \mathbf{Z}_{ijkl}^\top \mathbf{u} + \epsilon_{ijkl} \quad (1)$$

where i indexes items, j indexes studies, k indexes models, ℓ indexes training approaches, $\mathbf{X}$ contains fixed effects, $\boldsymbol{\beta}$ are fixed-effect coefficients, $\mathbf{Z}$ is the design matrix for random effects, $\mathbf{u}$ are random effect realizations, and $\epsilon_{ijkl} \sim \mathcal{N}(0, \sigma_\epsilon^2)$.

The random effects progress from simpler two-way structures (e.g., random intercepts for item and implementation) to more fully crossed structures including study/model/training and their interactions, and finally to *random item slopes within each higher-level factor*. This last structure (Model 6 in Table 1) operationalizes a key practical concern: a model (or training approach) may perform well on one item and poorly on another, and this item-by-implementation instability is itself part of the phenomenon being studied.

Bayesian Estimation The most conservative item-varying specification is high-dimensional relative to the number of studies and training regimes. To avoid overconfidence from boundary estimates and to stabilize inference under small- K clustering, we re-estimate the Model 6 specification in a Bayesian framework. The item-varying random effects explicitly model a phenomenon that is operationally critical in educational scoring: *a model can be reliable on some items and unreliable on others, and this pattern differs across studies and training regimes*. Details are in Appendix E.1

5 Results and Discussion

Across studies, we observe: (a) lower y_{ij} on reading versus science items, consistent with greater dependence on deep semantic integration; (b) decoder-only architectures lag encoder-based, aligning with masked-language-model pretraining advantages in bidirectional semantic representation; (c) tokenizer vocabulary size correlates with performance up to a point, after which gains taper, consistent with a concave $f(\cdot)$. Moreover, two practical constraints emerge: limited transferability of fine-tuned models across item sets and pronounced

sensitivity to prompt wording. Both phenomena indicate that models leverage superficial distributional cues rather than robust, rubric-aligned constructs. Findings in Table 2 are robust to both frequentist and Bayesian frameworks.

Meaning dependence reliably lowers agreement.

The read indicator is negative across specifications, including the Bayesian model (posterior mean ≈ -0.21 on Fisher- z scale). This implies that when an item requires integrating meaning rather than matching surface features, LLM-based scorers exhibit a systematic drop in agreement with humans. Practically, this is the regime where rubric interpretation, negation, partial credit logic, and child-specific phrasing all matter—and where next-token training objectives and instruction-following priors can yield plausible but rubric-inconsistent judgments.

Decoder (GPT-style) architectures are not uniformly advantaged for reliability.

The gpt coefficient is negative in the Bayesian specification (posterior mean ≈ -0.39 , with the 95% CI excluding 0 in the provided fit). This result links an architectural propensity—autoregressive generation optimized for fluent continuation—to a reliability-style evaluation criterion. Decoder models may be especially prone to (i) over-conditioning on spurious lexical cues, (ii) verbosity or rationale-driven confabulation that is not anchored to the rubric, and (iii) sensitivity to prompt framing that manifests as higher variance across items and studies. Even when average performance is acceptable, instability across items can remain substantial.

Tokenization matters, and vocabulary size exhibits a “Goldilocks” region.

Tokenization family coefficients are all strongly positive because the model is fit without an intercept; they represent baseline expected performance under each tokenizer when other covariates are at their reference levels. The more diagnostic result is the quadratic vocabulary effect: $|V|^2$ is negative (Bayes ≈ -0.06 , excluding 0), consistent with a concave-down relationship. This supports a practical hypothesis specific to child-written responses: extremely small vocabularies over-fragment idiosyncratic spellings and morphology, while extremely large vocabularies include rare or under-trained tokens that behave unpredictably on out-of-distribution orthography. Both failure modes reduce scoring reliability, even if they do not strongly

affect perceived fluency.

Scaling helps, but not as a primary lever

Model size (logsize) is positive in the best-controlled models (Bayes ≈ 0.06), indicating that scaling can improve agreement, but the effect is modest. This aligns with practitioner experience: larger LLMs may become better at paraphrasing and world knowledge, but rubric-aligned ordinal scoring remains bottlenecked by stability, calibration, and robustness to child-language noise.

Human item difficulty $\neq$ LLM difficulty A central and practically important null result is that QWK_{hum} (humqwk) is not reliably associated with LLM QWK in any parameterization, including the Bayesian model (posterior mean near 0 with wide uncertainty). This disconnect is informative:

- Items “hard for humans” (low human–human agreement) are not necessarily hard for LLMs.
- Conversely, items that humans score consistently can still be difficult for LLMs, particularly when they require semantic integration under noisy writing.

This motivates a methodological caution for deployment: practitioners should avoid anthropomorphizing item difficulty. Human rater disagreement reflects ambiguity, rubric underspecification, and subjective interpretation; LLM failure often reflects distribution shift (child orthography), tokenization artifacts, and sensitivity to prompt or training regimen. The two difficulty notions are not interchangeable, so human reliability metrics should be used as *quality control for labels*, not as a proxy for expected model behavior.

5.1 Argument from silence is shockingly loud.

While performing the systematic review and aggregations of data for the present study, we note that many published studies use modern LLMs for short answer scoring using the ASAP-SAS dataset (see Appendix C.1). Despite the accepted metric being QWK and despite there being 10 total prompts, nearly every paper (except those evaluated herein) either chose not to report QWK or chose not to report on all 10 prompts. This is not a coincidence. While not the intent of this study, it suggests that the most underreported SAS tasks are those associated with statistical artifacts of autoregression. Bluntly, studies are shielding null or weak results.

6 Implications and Recommendations

Design for robustness to child text. Tokenization and vocabulary design are not incidental implementation details; they are part of the reliability mechanism. When selecting or training models for child-written responses, practitioners should explicitly test sensitivity to misspellings, invented words, and nonstandard morphology, and avoid extremes of vocabulary size.

Treat semantic items as a separate regime.

Meaning-dependent items should be evaluated and monitored separately. If an assessment mixes surface-aligned and meaning-dependent items, a single overall QWK obscures the most consequential failures.

Recommendations Redirect effort from prompt engineering and in-context tricks toward: (i) item-level assessment-aligned training objectives that directly optimize agreement with rubric-anchored, human-verified labels; (ii) meaning-grounded representations, e.g., encoder backbones or dual encoders that explicitly model evidence extraction and claim–evidence alignment; (iii) uncertainty-aware scoring that reports calibrated confidence alongside QWK expectations; (iv) rigorous, mixed-effects benchmarking across item sets to guard against overfitting; (v) new datasets spanning contemporary curricula, to ensure generalization beyond legacy competitions; (vi) use of interpretable, interchangeable (see Kim et al.) automated scoring mechanisms where establishment of train.

7 Conclusion

Autoscoring remains the most important unsolved challenge in edtech because it requires trustworthy judgments about learning. Our meta-analysis clarifies why progress has stalled: autoregressive training steers models toward surface form, away from the construct validity autoscoring demands. Closing the gap will require models, objectives, and evaluations purpose-built for educational assessment. Practically, that means investing in validity studies, reporting score uncertainty, and designing pipelines where models justify scores with evidence traces. Scientifically, it means testing hypotheses like ours on newer corpora and releasing transparent benchmarks so progress reflects genuine improvements in educational measurement, not proxy leaderboard gains.

Table 2: Meta-regression Estimated Parameters

	(1)	(2)	(3)	(4)	(5)	(6)	(6, Bayes)
tok:BPE	0.75*** [0.47, 0.89]	0.78*** [0.52, 0.91]	0.80*** [0.41, 0.94]	0.81*** [0.51, 0.93]	0.80*** [0.51, 0.93]	0.72*** [0.63, 0.79]	0.78*** [0.50, 0.92]
tok:Unigram	0.66*** [0.33, 0.85]	0.73*** [0.44, 0.89]	0.76*** [0.33, 0.93]	0.77*** [0.44, 0.92]	0.77*** [0.44, 0.91]	0.71*** [0.63, 0.78]	0.76*** [0.44, 0.90]
tok:WordP	0.65*** [0.33, 0.84]	0.74*** [0.45, 0.89]	0.78*** [0.37, 0.93]	0.78*** [0.46, 0.92]	0.77*** [0.46, 0.92]	0.72*** [0.64, 0.79]	0.77*** [0.48, 0.91]
vocab	0.35** [0.13, 0.54]	-0.03 [-0.24, 0.18]	0.12 [-0.12, 0.35]	0.04 [-0.15, 0.23]	0.04 [-0.15, 0.22]	0.10 [-0.04, 0.24]	0.10 [-0.04, 0.24]
vocab ²	-0.12* [-0.22, -0.01]	0.00 [-0.08, 0.08]	-0.07 [-0.16, 0.03]	-0.03 [-0.10, 0.04]	-0.03 [-0.10, 0.04]	-0.06* [-0.12, -0.01]	-0.06* [-0.11, -0.00]
QWK _{hum}	0.03 [-0.21, 0.27]	0.03 [-0.22, 0.28]	0.03 [-0.22, 0.28]	0.02 [-0.20, 0.24]	0.03 [-0.19, 0.24]	-0.02 [-0.07, 0.04]	0.00 [-0.27, 0.28]
read item	-0.22** [-0.37, -0.05]	-0.22* [-0.37, -0.05]	-0.22* [-0.37, -0.05]	-0.21** [-0.35, -0.06]	-0.21** [-0.35, -0.07]	-0.23*** [-0.26, -0.21]	-0.21** [-0.37, -0.03]
gpt-style	-0.33*** [-0.49, -0.14]	-0.25** [-0.42, -0.07]	-0.46+ [-0.78, 0.05]	-0.40+ [-0.70, 0.03]	-0.40+ [-0.70, 0.03]	-0.14 [-0.31, 0.04]	-0.37*** [-0.58, -0.14]
logsize	-0.07** [-0.12, -0.02]	0.01 [-0.05, 0.07]	0.07* [0.01, 0.13]	0.05* [0.01, 0.10]	0.05* [0.01, 0.10]	0.06*** [0.03, 0.09]	0.06** [0.03, 0.10]
<i>Rand. Effects</i>							
LLM instance	Y						
LLM		Y	Y	Y	Y	Y	Y
Training			Y	Y	Y	Y	Y
Study		Y	Y	Y	Y	Y	Y
Item	Y	Y	Y	Y	Y		
2° Interactions				Y	Y		
3° Interactions					Y		
<i>Rand. Slopes</i>							
Items						Y	Y
N	890	890	890	890	890	890	890
R^2_{Marg}	0.456	0.240	0.227	0.230	0.235	0.143	0.235
R^2_{Cond}	0.709	0.609	0.823	0.876	0.876	0.898	0.875
AIC	-322.5	-381.3	-680.5	-939.6	-934.0	-797.9	
BIC	-265.0	-319.0	-613.4	-843.8	-823.8	40.6	
RMSE/ σ_ϵ	0.17	0.18	0.15	0.10	0.10	0.10	0.10/0.12

We report fixed-effect estimates with 95% uncertainty intervals. Because effects are estimated on Fisher- z scale, coefficients are mapped back to the QWK scale: $\hat{\kappa}_w = \tanh(\hat{\eta})$. **Model 1** included random intercepts for items and implementations (unique combinations of model, training approach, and study): $\mathbf{u} = \{u_{\text{item}_i}, u_{\text{implementation}_{jkl}}\}$. **Model 2** decomposed implementation variance into study and model components: $\mathbf{u} = \{u_{\text{item}_i}, u_{\text{study}_j}, u_{\text{model}_k}\}$. **Model 3** further separated training approach effects: $\mathbf{u} = \{u_{\text{item}_i}, u_{\text{study}_j}, u_{\text{model}_k}, u_{\text{training}_\ell}\}$. **Model 4** added all two-way interactions among item, study, model, and training: $\mathbf{u} = \{u_{\text{item}_i}, u_{\text{study}_j}, u_{\text{model}_k}, u_{\text{training}_\ell}, u_{\text{item} \times \text{study}_{ij}}, u_{\text{item} \times \text{model}_{ik}}, \dots\}$. **Model 5** extended to selected three-way interactions of theoretical interest. **Model 6** employed random slopes, allowing item effects to vary across study, model, and training contexts: $\mathbf{u} = \{u_{\text{study}_j} \sim \mathcal{N}(\mathbf{0}, \Sigma_{\text{study}}), u_{\text{model}_k} \sim \mathcal{N}(\mathbf{0}, \Sigma_{\text{model}}), u_{\text{training}_\ell} \sim \mathcal{N}(\mathbf{0}, \Sigma_{\text{training}})\}$, where each $\mathbf{u}$ is a 10-dimensional vector (one component per item) with full covariance matrix Σ . *** < 0.005 , ** < 0.01 , * < 0.05 , + < 0.1

Thus, autoscoring is not merely another “apply LLMs with better prompts” problem. The gap to human raters persists because the optimization target is misaligned with the assessment target (McCoy et al., 2023; Jurenka et al., 2024). Second, the reading-versus-science disparity cautions against averaging scores across domains; meaning dependence matters. Third, architecture and pretraining matter: encoder or hybrid systems with objectives that force bidirectional semantic integration are better suited than pure decoders for rubric-grounded judgments. Fourth, tokenization choices influence performance, but they are not a panacea; beyond a threshold, scaling vocabulary does not fix construct underrepresentation.

This research demonstrates that the path to effective autoscoring does not lie in simply scaling up existing LLMs or refining prompts. The anticlimactic performance of these models highlights a fundamental limitation that must be addressed at a deeper level. Future work in this critical area of educational technology should pivot away from these superficial approaches. Instead, the field must focus on developing models and training regimes specifically designed for the challenges of educa-

tional assessment. This requires a concerted effort, combining deep expertise in both language models and educational measurement, to build systems that can robustly and fairly answer the most important question: is a child learning?

8 Limitations and future work

First, QWK—while appropriate for ordinal agreement—still compresses rich error structure into a single coefficient; future work should meta-analyze score-category confusion patterns or expected quadratic loss directly. Second, vocabulary size and tokenizer family are partially confounded with model families in the ecosystem; causal interpretations should be made cautiously, even though the multiway random effects substantially reduce spurious certainty. Only the direction (and not the magnitude) of the estimated coefficients should be considered. Third, we analyzed Fisher- z transformed QWK with a Gaussian likelihood; a more explicit measurement model for kappa-type statistics could further improve calibration when effects approach the boundaries, especially of patterns found in effect-sizes that represent underlying imbalanced datasets.

Despite these limitations, the convergence of results across diverse specifications—including a Bayesian fit designed for small- K , high-heterogeneity settings—supports a coherent conclusion: for scoring child-written short answers, reliability of humanlike scoring is governed less by human-perceived item difficulty and more by architecture–tokenization interactions and semantic dependence, with substantial multi-level instability that must be modeled rather than averaged away.

References

- Muhammad Abbas, Farooq Ahmed Jam, and Tariq Iqbal Khan. 2024. [Is it harmful or helpful? Examining the causes and consequences of generative AI usage among university students](#). *International Journal of Educational Technology in Higher Education*, 21(1):10.
- Khan Academy. 2023. [Khan Academy’s approach to the responsible development of AI - Khan Academy Blog](#).
- April Baker-Bell. 2020. *Linguistic Justice: Black Language, Literacy, Identity, and Pedagogy*. Routledge, New York.
- Solon Barocas, Anhong Guo, Ece Kamar, Jacquelyn Krones, Meredith Ringel Morris, Jennifer Wortman Vaughan, W. Duncan Wadsworth, and Hanna Wallach. 2021. [Designing Disaggregated Evaluations of AI Systems: Choices, Considerations, and Tradeoffs](#). In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, AIES ’21, pages 368–378, New York, NY, USA. Association for Computing Machinery.
- Douglas Bates, Martin Mächler, Ben Bolker, and Steve Walker. 2015. [Fitting Linear Mixed-Effects Models Using lme4](#). *Journal of Statistical Software*, 67:1–48.
- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. [On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?](#) In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT ’21, pages 610–623, New York, NY, USA. Association for Computing Machinery.
- Ruha Benjamin. 2019. *Race after technology: abolitionist tools for the New Jim Code*. Polity, Cambridge, UK Medford, MA.
- James Bisbee, Joshua D. Clinton, Cassy Dorff, Brenton Kenkel, and Jennifer M. Larson. 2024. [Synthetic Replacements for Human Survey Data? The Perils of Large Language Models](#). *Political Analysis*, pages 1–16.
- Hezekiah J. Branch, Jonathan Rodriguez Cefalu, Jeremy McHugh, Leyla Hujer, Aditya Bahl, Daniel del Castillo Iglesias, Ron Heichman, and Ramesh Darwishi. 2022. [Evaluating the Susceptibility of Pre-Trained Language Models via Handcrafted Adversarial Examples](#). *arXiv preprint*. ArXiv:2209.02128.
- Tanise Ceron, Neele Falk, Ana Barić, Dmitry Nikolaev, and Sebastian Padó. 2024. [Beyond prompt brittleness: Evaluating the reliability and consistency of political worldviews in LLMs](#). *arXiv preprint*. ArXiv:2402.17649.
- Imran Chamieh, Torsten Zesch, and Klaus Giebertmann. 2024. [LLMs in Short Answer Scoring: Limitations and Promise of Zero-Shot and Few-Shot Approaches](#). In *Proceedings of the 19th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2024)*, pages 309–315, Mexico City, Mexico. Association for Computational Linguistics.
- Myra Cheng, Esin Durmus, and Dan Jurafsky. 2023. [Marked Personas: Using Natural Language Prompts to Measure Stereotypes in Language Models](#). *arXiv preprint*. ArXiv:2305.18189.
- Julian Coda-Forno, Kristin Witte, Akshay K. Jagadish, Marcel Binz, Zeynep Akata, and Eric Schulz. 2023. [Inducing anxiety in large language models increases exploration and bias](#). *arXiv preprint*. ArXiv:2304.11111.
- Nicholas Deas, Jessica Grieser, Shana Kleiner, Desmond Patton, Elsbeth Turcan, and Kathleen McKeown. 2023. [Evaluation of African American Language Bias in Natural Language Generation](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 6805–6824, Singapore. Association for Computational Linguistics.
- Vittoria Dentella, Fritz Günther, and Evelina Leivada. 2023. [Systematic testing of three Language Models reveals low language accuracy, absence of response stability, and a yes-response bias](#). *Proceedings of the National Academy of Sciences*, 120(51):e2309583120. Publisher: Proceedings of the National Academy of Sciences.
- Ameet Deshpande, Vishvak Murahari, Tanmay Rajpurohit, Ashwin Kalyan, and Karthik Narasimhan. 2023. [Toxicity in chatgpt: Analyzing persona-assigned language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 1236–1270, Singapore. Association for Computational Linguistics.
- 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, Minneapolis, Minnesota. Association for Computational Linguistics.

- Allyson Ettinger. 2020. [What BERT Is Not: Lessons from a New Suite of Psycholinguistic Diagnostics for Language Models](#). *Transactions of the Association for Computational Linguistics*, 8:34–48. Place: Cambridge, MA.
- Philip Gage. 1994. A new algorithm for data compression. *C Users J.*, 12(2):23–38.
- Jonas Geiping, Alex Stein, Manli Shu, Khalid Saifullah, Yuxin Wen, and Tom Goldstein. 2024. [Coercing LLMs to do and reveal \(almost\) anything](#). *arXiv preprint*. ArXiv:2402.14020 [cs].
- Hila Gonen and Yoav Goldberg. 2019. [Lipstick on a Pig: Debiasing Methods Cover up Systematic Gender Biases in Word Embeddings But do not Remove Them](#). 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 609–614, Minneapolis, Minnesota. Association for Computational Linguistics.
- Kai Greshake, Sahar Abdelnabi, Shailesh Mishra, Christoph Endres, Thorsten Holz, and Mario Fritz. 2023. [Not what you’ve signed up for: Compromising Real-World LLM-Integrated Applications with Indirect Prompt Injection](#).
- Shashank Gupta, Vaishnavi Shrivastava, Ameet Deshpande, Ashwin Kalyan, Peter Clark, Ashish Sabharwal, and Tushar Khot. 2024. [Bias Runs Deep: Implicit Reasoning Biases in Persona-Assigned LLMs](#). *arXiv preprint*. ArXiv:2311.04892.
- Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel Bowman, and Noah A. Smith. 2018. [Annotation Artifacts in Natural Language Inference Data](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 107–112, New Orleans, Louisiana. Association for Computational Linguistics.
- Michael Hardy. 2024. Rater Effects vs Robo-Effects: Profiling Automated Classroom Evaluation. Writing Assessments and Rater Effects, Philadelphia, PA.
- Michael Hardy. 2025. [“All that Glitters”: Techniques for Evaluations with Unreliable Model and Human Annotations](#). In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 2250–2278, Albuquerque, New Mexico. Association for Computational Linguistics.
- Valentin Hofmann, Pratyusha Ria Kalluri, Dan Jurafsky, and Sharese King. 2024a. [AI generates covertly racist decisions about people based on their dialect](#). *Nature*, pages 1–8.
- Valentin Hofmann, Pratyusha Ria Kalluri, Dan Jurafsky, and Sharese King. 2024b. [Dialect prejudice predicts AI decisions about people’s character, employability, and criminality](#). *arXiv preprint*. ArXiv:2403.00742 [cs].
- 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*, pages 438–442, New York, NY, USA. Association for Computing Machinery.
- Irina Jurenka, Markus Kunesch, Kevin R McKee, Daniel Gillick, Shaojian Zhu, Shubham Milind Phal, Katherine Hermann, Daniel Kasenberg, Avishkar Bhoopchand, Ankit Anand, Miruna Pislari, Stephanie Chan, Lisa Wang, Jennifer She, Parsa Mahmoudieh, Wei-Jen Ko, Andrea Huber, Brett Wiltshire, Gal Elidan, and 51 others. 2024. Towards Responsible Development of Generative AI for Education: An Evaluation-Driven Approach.
- Nora Kassner and Hinrich Schütze. 2020. [Negated and Misprimed Probes for Pretrained Language Models: Birds Can Talk, But Cannot Fly](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7811–7818, Online. Association for Computational Linguistics.
- Douwe Kiela, Tristan Thrush, and Kawin Ethayarajh. 2023. [Plotting Progress in AI](#).
- Najoung Kim and Sebastian Schuster. 2023. [Entity Tracking in Language Models](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3835–3855, Toronto, Canada. Association for Computational Linguistics.
- Yunsung Kim, Mike Hardy, Joseph Tey, Candace Thille, and Chris Piech. 2026. [Principled Design of Interpretable Automated Scoring for Large-Scale Educational Assessments](#). *arXiv preprint*. ArXiv:2511.17069 [cs] version: 2.
- David Klahr. 2013. [What do we mean? On the importance of not abandoning scientific rigor when talking about science education](#). *Proceedings of the National Academy of Sciences*, 110(supplement_3):14075–14080.
- Artur Klingbeil, Cassandra Grützner, and Philipp Schreck. 2024. [Trust and reliance on AI — An experimental study on the extent and costs of overreliance on AI](#). *Computers in Human Behavior*, 160:108352.
- Ryan Koo, Minhwa Lee, Vipul Raheja, Jong Inn Park, Zae Myung Kim, and Dongyeop Kang. 2024. [Benchmarking Cognitive Biases in Large Language Models as Evaluators](#). In *Findings of the Association for Computational Linguistics ACL 2024*, pages 517–545, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.
- Eliza Kosoy, Soojin Jeong, Anoop Sinha, Alison Gopnik, and Tanya Kraljic. 2024. [Children’s Mental Models of Generative Visual and Text Based AI Models](#).
- Eliza Kosoy, Emily Rose Reagan, Leslie Lai, Alison Gopnik, and Danielle Krettek Cobb. 2023. [Comparing Machines and Children: Using Developmental](#)

- Psychology Experiments to Assess the Strengths and Weaknesses of LaMDA Responses. *arXiv preprint*. ArXiv:2305.11243 [cs].
- Hadas Kotek, Rikker Dockum, and David Sun. 2023. Gender bias and stereotypes in Large Language Models. In *Proceedings of The ACM Collective Intelligence Conference, CI '23*, pages 12–24, New York, NY, USA. Association for Computing Machinery.
- Taku Kudo. 2018. Subword Regularization: Improving Neural Network Translation Models with Multiple Subword Candidates. *arXiv preprint*. ArXiv:1804.10959 [cs] version: 1.
- Taku Kudo and John Richardson. 2018. SentencePiece: A simple and language independent subword tokenizer and detokenizer for Neural Text Processing. *arXiv preprint*. ArXiv:1808.06226 [cs].
- Sander Land and Max Bartolo. 2024. Fishing for Magikarp: Automatically Detecting Under-trained Tokens in Large Language Models. *arXiv preprint*. ArXiv:2405.05417 [cs].
- Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, Benjamin Newman, Binhang Yuan, Bobby Yan, Ce Zhang, Christian Cosgrove, Christopher D. Manning, Christopher Ré, Diana Acosta-Navas, Drew A. Hudson, and 31 others. 2023. Holistic Evaluation of Language Models. *arXiv preprint*. ArXiv:2211.09110 [cs].
- Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2023. Lost in the Middle: How Language Models Use Long Contexts. *arXiv preprint*. ArXiv:2307.03172 [cs].
- Nelson F. Liu, Roy Schwartz, and Noah A. Smith. 2019. Inoculation by Fine-Tuning: A Method for Analyzing Challenge Datasets. *arXiv preprint*. ArXiv:1904.02668 [cs].
- Qi Liu, Yan Zhuang, Haoyang Bi, Zhenya Huang, Weizhe Huang, Jiatong Li, Junhao Yu, Zirui Liu, Zirui Hu, Yuting Hong, Zachary A. Pardos, Haiping Ma, Mengxiao Zhu, Shijin Wang, and Enhong Chen. 2024. Survey of Computerized Adaptive Testing: A Machine Learning Perspective. *arXiv preprint*. ArXiv:2404.00712 version: 2.
- Kyle Mahowald, Anna A. Ivanova, Idan A. Blank, Nancy Kanwisher, Joshua B. Tenenbaum, and Evelina Fedorenko. 2023. Dissociating language and thought in large language models.
- R. Thomas McCoy, Shunyu Yao, Dan Friedman, Matthew Hardy, and Thomas L. Griffiths. 2023. Embers of Autoregression: Understanding Large Language Models Through the Problem They are Trained to Solve.
- Sewon Min, Xinxu Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2022. Rethinking the Role of Demonstrations: What Makes In-Context Learning Work? *arXiv preprint*. ArXiv:2202.12837 [cs].
- Huy Nghiem, John Prindle, Jieyu Zhao, and Hal Daumé III. 2024. "You Gotta be a Doctor, Lin": An Investigation of Name-Based Bias of Large Language Models in Employment Recommendations. *arXiv preprint*. ArXiv:2406.12232.
- Allen Nie, Yash Chandak, Miroslav Suzara, Malika Ali, Juliette Woodrow, Matt Peng, Mehran Sahami, Emma Brunskill, and Chris Piech. 2024. The GPT Surprise: Offering Large Language Model Chat in a Massive Coding Class Reduced Engagement but Increased Adopters Exam Performances. *arXiv preprint*. ArXiv:2407.09975 [cs, stat].
- Open Innovation Team and Department for Education. 2024. Generative AI in Education: Educator and Expert Views. Technical report, UK Department for Education. ERIC Number: ED649949.
- OpenAI. 2024. openai/tiktoken. Original-date: 2022-12-01T23:22:11Z.
- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, and 262 others. 2023. GPT-4 Technical Report.
- Christopher Michael Ormerod. 2022. Short-answer scoring with ensembles of pretrained language models. *arXiv preprint*. ArXiv:2202.11558 [cs].
- Christopher Michael Ormerod and Alexander Kwako. 2024. Automated Text Scoring in the Age of Generative AI for the GPU-poor. *arXiv preprint*. ArXiv:2407.01873 [cs] version: 1.
- Richard Yuanzhe Pang, Vishakh Padmakumar, Thibault Sellam, Ankur P. Parikh, and He He. 2023. Reward Gaming in Conditional Text Generation. *arXiv preprint*. ArXiv:2211.08714 [cs].
- Martin Pawelczyk, Seth Neel, and Himabindu Lakkaraju. 2024. In-Context Unlearning: Language Models as Few Shot Unlearners. *arXiv preprint*. ArXiv:2310.07579.
- Adam Poliak, Jason Naradowsky, Aparajita Haldar, Rachel Rudinger, and Benjamin Van Durme. 2018. Hypothesis Only Baselines in Natural Language Inference. In *Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics*, pages 180–191, New Orleans, Louisiana. Association for Computational Linguistics.
- Michael J D Powell. 2009. The BOBYQA algorithm for bound constrained optimization without derivatives. *Cambridge NA Report*, NA2009/06:26–46.

- Sippo Rossi, Alisia Marianne Michel, Raghava Rao Mukkamala, and Jason Bennett Thatcher. 2024. [An Early Categorization of Prompt Injection Attacks on Large Language Models](#). *arXiv preprint*. ArXiv:2402.00898 [cs].
- Sander Schulhoff, Michael Ilie, Nishant Balepur, Konstantine Kahadze, Amanda Liu, Chenglei Si, Yinheng Li, Aayush Gupta, Hyojung Han, Sevien Schulhoff, Pranav Sandeep Dulepet, Saurav Vidyadhara, Dayeon Ki, Sweta Agrawal, Chau Pham, Gerson Kroiz, Feileen Li, Hudson Tao, Ashay Srivastava, and 12 others. 2024. [The Prompt Report: A Systematic Survey of Prompting Techniques](#).
- Sander Schulhoff, Jeremy Pinto, Anaum Khan, Louis-François Bouchard, Chenglei Si, Svetlana Anati, Valen Tagliabue, Anson Liu Kost, Christopher Carnahan, and Jordan Boyd-Graber. 2023. [Ignore This Title and HackAPrompt: Exposing Systemic Vulnerabilities of LLMs through a Global Scale Prompt Hacking Competition](#).
- Melanie Sclar, Yejin Choi, Yulia Tsvetkov, and Alane Suhr. 2024. [Quantifying Language Models’ Sensitivity to Spurious Features in Prompt Design or: How I learned to start worrying about prompt formatting](#). *arXiv preprint*. ArXiv:2310.11324 [cs].
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2015. [Neural Machine Translation of Rare Words with Subword Units](#).
- Mark D. Shermis. 2015. [Contrasting State-of-the-Art in the Machine Scoring of Short-Form Constructed Responses](#). *Educational Assessment*, 20(1):46–65. Publisher: Routledge _eprint: <https://doi.org/10.1080/10627197.2015.997617>.
- Mark D. Shermis and Jill Burstein, editors. 2003. *Automated essay scoring: a cross-disciplinary perspective*. L. Erlbaum Associates, Mahwah, N.J.
- Freda Shi, Xinyun Chen, Kanishka Misra, Nathan Scales, David Dohan, Ed Chi, Nathanael Schärli, and Denny Zhou. 2023. [Large Language Models Can Be Easily Distracted by Irrelevant Context](#). *arXiv preprint*. ArXiv:2302.00093 [cs].
- Evan Shieh, Faye-Marie Vassel, Cassidy Sugimoto, and Thema Monroe-White. 2024. [Laissez-Faire Harms: Algorithmic Biases in Generative Language Models](#). *arXiv preprint*. ArXiv:2404.07475.
- Mithunan Sivakumar, Ali Shariq Imran, Zenun Kastrati, Ahmet Soylu, and Muhamet Kastrati. 2026. [The Future of Grading: Can LLMs Accurately Score Student Work?](#) In *Proceedings of the Future Technologies Conference (FTC) 2025, Volume 4*, pages 268–289, Cham. Springer Nature Switzerland.
- R Core Team. [R: A Language and Environment for Statistical Computing](#).
- Lindia Tjuatja, Valerie Chen, Sherry Tongshuang Wu, Ameet Talwalkar, and Graham Neubig. 2024. [Do LLMs exhibit human-like response biases? A case study in survey design](#). *arXiv preprint*. ArXiv:2311.04076.
- Aleksandar Todorov, Elisa Klunder, and Julia Belloni. 2025. [Evaluating the Potential of LLMs for Better Short Answer Scoring](#). In *Proceedings of the 17th International Conference on Computer Supported Education*, pages 108–119, Porto, Portugal. SCITEPRESS - Science and Technology Publications.
- Angelina Wang, Jamie Morgenstern, and John P. Dickerson. 2024a. [Large language models should not replace human participants because they can misportray and flatten identity groups](#). *arXiv preprint*. ArXiv:2402.01908.
- Dixuan Wang, Yanda Li, Junyuan Jiang, Zepeng Ding, Guochao Jiang, Jiaqing Liang, and Deqing Yang. 2024b. [Tokenization Matters! Degrading Large Language Models through Challenging Their Tokenization](#).
- Melissa Warr. 2024. [Beat Bias? Personalization, Bias, and Generative AI](#). pages 1481–1488. Association for the Advancement of Computing in Education (AACE).
- Melissa Warr, Nicole Jakubczyk Oster, and Roger Isaac. 2024. [Implicit bias in large language models: Experimental proof and implications for education](#). *Journal of Research on Technology in Education*, 0(0):1–24. _eprint: <https://doi.org/10.1080/15391523.2024.2395295>.
- Albert Webson and Ellie Pavlick. 2022. [Do Prompt-Based Models Really Understand the Meaning of Their Prompts?](#) In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2300–2344, Seattle, United States. Association for Computational Linguistics.
- Jiaxin Wen, Ruiqi Zhong, Akbir Khan, Ethan Perez, Jacob Steinhardt, Minlie Huang, Samuel R. Bowman, He He, and Shi Feng. 2024. [Language Models Learn to Mislead Humans via RLHF](#).
- Congying Xia, Chen Xing, Jiangshu Du, Xinyi Yang, Yihao Feng, Ran Xu, Wenpeng Yin, and Caiming Xiong. 2024. [FOFO: A Benchmark to Evaluate LLMs’ Format-Following Capability](#). *arXiv preprint*. ArXiv:2402.18667.
- Changrong Xiao, Wenxing Ma, Sean Xin Xu, Kunpeng Zhang, Yufang Wang, and Qi Fu. 2024. [From Automation to Augmentation: Large Language Models Elevating Essay Scoring Landscape](#).
- Jintang Xue, Yun-Cheng Wang, Chengwei Wei, Xiaofeng Liu, Jonghye Woo, and C.-C. Jay Kuo. 2023. [Bias and Fairness in Chatbots: An Overview](#). *arXiv preprint*. ArXiv:2309.08836.

Eunice Yiu, Eliza Kosoy, and Alison Gopnik. 2023. [Imitation versus Innovation: What children can do that large language and language-and-vision models cannot \(yet\)?](#) *arXiv preprint*. ArXiv:2305.07666 [cs].

Chunpeng Zhai, Santoso Wibowo, and Lily D. Li. 2024. [The effects of over-reliance on AI dialogue systems on students’ cognitive abilities: a systematic review](#). *Smart Learning Environments*, 11(1):28.

Kaitlyn Zhou, Jena D. Hwang, Xiang Ren, and Maarten Sap. 2024. [Relying on the Unreliable: The Impact of Language Models’ Reluctance to Express Uncertainty](#). *arXiv preprint*. ArXiv:2401.06730 [cs].

Kaitlyn Zhou, Dan Jurafsky, and Tatsunori Hashimoto. 2023. [Navigating the Grey Area: How Expressions of Uncertainty and Overconfidence Affect Language Models](#). *arXiv preprint*. ArXiv:2302.13439 [cs].

A Token Sensitivity Experiment

In this analysis we use an edtech industry-leader’s teacher assistant tool to generate documentation for special education students that is measurably racially biased in a way that would exacerbate existing issues of inequity. The results of that study can be seen in Table 4 and full dataset is available online.⁴ **We conjecture that increased presence of subsequent wording and tokenization issues could amplify these effects further, as has been seen in other areas of research** (Cheng et al., 2023; Kotek et al., 2023; Wang et al., 2024a).

Khanmigo,⁵ the Generative AI assistant from Khan Academy, offers various conversational teacher supports, including an Individualized Education Program (IEP) Assistant.⁶ IEPs are an important use case because they are student-specific, and in the United States, legally binding documents required for special education children that outlines the intended educational experience to be provided. In a sense, understanding these documents are a critical step towards individualized learning. But more importantly, these tasks are difficult to measure for quality, as each is intended to be unique to each child.

Khan Academy uses fine-tuning, prompt engineering, “red teaming”, and monitoring for inappropriate user interactions to mitigate risks and bias (Academy, 2023), illustrating many current best practices for responsible AI and for the reduction

of biases for educational use cases. This study therefore asks whether current practices in edtech can sufficiently minimize potential harms to children arising from propagated racial biases using a disaggregated evaluation (Barocas et al., 2021).

Table 3: IEP assistant input set-up

	Prompt Slots	Experimental Input
(a)	Learning environment	“third grade math class”
(b)	Disability category	“student is ****”
(c)	Current challenges	“struggling in class”
(d)	Current strengths	“good at sports”

The Khanmigo IEP Assistant interface has four prompt slots for providing information about the child for whom the IEP is being generated, as seen in Table 3. The interface used and descriptions of categories can be found in Appendix Section A.1. Each race/ethnicity replaced the “****” in prompt slot (b).

We exploit known GPT architecture weaknesses, including sensitivities to tokenization and formatting (Liang et al., 2023; Sclar et al., 2024; Wang et al., 2024b; Xia et al., 2024), as well as vulnerabilities to typos and irregular text (Greshake et al., 2023; Schulhoff et al., 2023; Sclar et al., 2024; Ceron et al., 2024; Tjautja et al., 2024; Webson and Pavlick, 2022).

Sensitivity to tokenization allows us to create many unique IEPs without any changes to text⁷ These sensitivities allow for the creation of unique IEPs through subtle manipulations of whitespace in prompts (b), (c), and (d), with many distinct permutations for each racial/ethnic group (Black, Hispanic, and White). For manual coding purposes, we generated 27 IEPs per group between July 20-30, 2024. For manual coding purposes, we only generate 27 for each student group.

After removing duplicates,⁸ 69 unique IEPs were analyzed and hand-coded to identify real-world differences in proposed interventions across demographics (Table 4). shows the distribution of the AI generated IEP educational interventions for children across demographics. While a quantitative evaluation using n-gram TF-IDF features could

⁷At the time of the experiments, we suspect that the temperature of this model was at or near zero. If the same IEP input was submitted, without changes to whitespace, an identical IEP was returned. Adding of whitespace was sufficient permutation for the diversity of responses.

⁸There were also six IEPs removed because they ignored the 3rd grade math prompt and produced generic IEPs for 6th grade reading. These 6th grade IEPs also resulted in biased behaviors.

⁴Token sensitivity model output.

⁵<https://www.khanacademy.org/teacher/khanmigo-tools/>

⁶<https://www.khanacademy.org/teacher/khanmigo-tools/iep-assistant>

Table 4: AI Generated IEP Interventions by Student Race/Ethnicity

Category	Adjustments/Interventions	Black N=25	Hispanic N=23	White N=21	p-val*
Social Development	Group or team class activities	100%	100%	100%	
	Structured peer interactions	84%	70%	95%	0.082
	Student mentorship	8%	0%	24%	0.025
	Class leadership opportunities	40%	57%	24%	0.093
Math Content Remediation	Small group or 1-on-1 instruction (teacher)	28%	43%	81%	0.001
	Tutoring (not teacher)	4%	0%	29%	0.003
	Kinesthetic/movement-based math activities	100%	100%	100%	
Engagement with Math Lessons	Integrating sports-themed content	64%	78%	43%	0.049
	Guest speakers from sports professions	0%	17%	5%	0.017
	Creating games or opportunities for play	48%	35%	76%	0.017
	Coordinate with PE staff for math-PE activities	24%	39%	62%	0.032
Reducing Student Misbehavior	Use timers to help student stay on task	48%	4%	10%	<0.001
	Set short-term behavioral goals with student	68%	61%	24%	0.007
	Regular behavior check-ins to discuss progress	40%	39%	29%	0.7
	System for reinforcement of positive behaviors	100%	100%	100%	
	Consistent use of routines	60%	30%	5%	<0.001
	Allow movement/energy breaks during class	64%	83%	38%	0.009
Staff†	Assigned seating	28%	9%	5%	0.075
	Professional development for school staff	68%	57%	86%	0.11

*Fisher's exact test

†Including staff professional development in a student's IEP is not appropriate.

Proportion of IEPs with the given intervention for demographic with greatest value where with $p < 0.1$. Any intervention appearing less than 4 times in IEPs for any race/ethnicity are not included in analysis for practical and statistical simplicity. Highlight colors are for emphasis.

easily produce statistically significant (and racially biased) results (Shieh et al., 2024), our focus was on meaningful real-world differences in children's educational experiences were the generated intervention plans put to use.

These results show an alarming display of differences in treatment of children that would lead to exacerbation of inequity based on irrelevant demographic information. The most resource-intensive interventions (in the Math Content Remediation category of Table 4) were disproportionately recommended for White children. Conversely, Black children received more interventions aimed at curbing noncompliant behaviors.

A.1 Interface

The interface for the IEP assistant, including descriptions of each of the four prompt components is shown in Figure 2. Generated IEP data and coding for this experiment can be found online. All outputs for this study were generated by the IEP Assistant from 20 July to 30 July 2024.

Figure 2: Khanmigo IEP Assistant Entry Page

Troublingly, these discriminatory differences in the treatment of children, if implemented, would likely go undetected by IEP Assistant users. Each individual IEP generated is written with the same formatting and confident tone that practitioners may not suspect that the recommendations are any-

thing less than ideal. (Zhou et al., 2024, 2023) Only through the generation of many examples do such patterns emerge. Our findings suggest that, similar to how dialect (Hofmann et al., 2024b), name (Nghiem et al., 2024), and preferences (Warr, 2024) can activate harmful generated content arising from spurious correlations (Hofmann et al., 2024a), we could reasonably assume that other biases would arise from subtle differences in token arrangements or wording (Coda-Forno et al., 2023; Warr et al., 2024).

We contrast the above findings with hallucinations made by the IEP assistant. Most of the IEPs generated across demographic lines contained content inappropriate for actual implementation, regardless of input quality. The individualized learning program of a child is not a place to plan professional learning for staff, the final category of Table 4.⁹ This underscores the insufficient understanding of the GPT model of IEP implementation, which occasionally results in confusing or misplaced recommendations. Although such hallucinations should be more obvious to a trained expert, the subtle subversion of some students based on race is harder to detect in a single generation.¹⁰

Evaluation of efficacy and appropriateness of the generated IEP interventions for the hypothetical child are not explicitly reported since this experiment intentionally demonstrates poor user inputs—inputs with results that align to biases found in model persona stereotyping (Gupta et al., 2024). Nevertheless, it is safe to say that the GPT model used clearly does not have sufficient understanding of how such IEPs are implemented, sometimes leading to humorously confused recommendations often using educational jargon (e.g., requiring teachers to ensure “stretch breaks” are “directly related to lesson objectives”).

These findings do not imply that Khan Academy offers a subpar edtech product. Rather, they highlight the insufficiency of industry best practices in mitigating harms arising from biases and spurious correlations in GPT models. This study serves as a call to action for improved safeguards and more ro-

⁹While professional development may be needed, an IEP is an inadvisable place to document that need.

¹⁰One possible reaction to the experiment could be to suggest not using the language of race/ethnicity or modify the prompt to not provoke the model. This is countered with the rhetorical questions, “What are the combinations of words that will avoid this?” “What prompts will guarantee that no biases will arise?” “Demonstration of harm should not require finding creative ways to elicit such behaviors as a test for whether such harms are real

bust evaluation methods in AI-assisted educational planning.

B Main Study Data

This section provides overall details of the meta-dataset construction.

B.1 ASAP Short Answer Scoring Dataset

Despite sophisticated prompt engineering and fine-tuning and despite these data being part of a Kaggle competition from over a decade ago, state-of-the-art models, including frontier decoders, fall short of human-level agreement.

Table 5: ASAP-SAS dataset detail by item.

	Token Len.	Train	Valid	Test	Assessment Area
Q1	47.5 ± 22.2	1,341	331	557	Science
Q2	59.2 ± 22.6	1,024	254	426	
Q3	47.9 ± 14.6	1,445	363	406	Reading (Informational Text)
Q4	40.3 ± 15.5	1,308	349	295	
Q5	25.1 ± 21.5	1,459	336	598	Science
Q6	23.8 ± 22.6	1,418	379	599	
Q7	41.3 ± 25.1	1,432	367	599	Reading (Literature)
Q8	53.0 ± 32.6	1,446	353	599	
Q9	49.7 ± 36.3	1,453	345	599	Reading (Informational Text)
Q10	41.1 ± 28.5	1,314	326	546	Science

B.1.1 Reading Item Example

In depth analysis of the grading task in the Short Answer Scoring dataset with the greatest discrepancies between human, SOTA model, and GPT4 model performance —0.968, 0.734, and 0.495, respectively—highlights how current models may not see improved performance on new autoscoring tasks even with retrieval of curricular materials to support model scoring and merits further examination. This particular item elicits a greater than average proportion of LLM failure modes in SAS, based on the discrepancies in distributional performance vis-à-vis human experts shown in Figure 3.

Item Content Item 7 contains a reading passage (876 words) and the following question:

Identify ONE trait that can describe Rose based on her conversations with Anna or Aunt Kolab. Include ONE detail from the story that supports your answer.




Figure 3: QWK distribution of LLMs on ASAP-SAS dataset for meta-analysis models relative to human performance

The average response length is 46 words (see Table 5). These are the instructions of the 0-2 range scoring rubric:

This item requires a two-part response. The student must identify one character trait of Rose, as well as include story details that adequately support the answer. The response should be related to the conversations with Anna or Aunt Kolab. In order to receive ANY credit, the student must provide a creditable response to the first part of the question. An additional point is awarded if the student provides a creditable response to the second part of the question.

Illustration of Aspects of Semantic Meaning Difficult for LLMs Through mini replications of the various experiments found across the papers in this study, we found that a common failure mode for this item was an inability for the model to distinguish between a “character trait”, a more enduring attribute of characters in the story that move beyond the scene they are in, and common uses of the word “trait” which may be interpreted to be any descriptor of an entity. When children would erroneously write, for example, that Rose’s character is

“frustrated” or “helping” or “caring for her sister”, LLMs often failed to recognize that these were not character traits, but were just actions or emotions observed during the scene. While this is easy for humans to do (having the highest human-human QWK), this meaning-based item was the hardest for LLMs. Anecdotally, even as we supplied LLMs with more detailed instructions about this particular point, the entire reading passage, and reasoning prompting, we did not see much change in LLM abilities to be successful on this meaning-based task.

C Meta-study Dataset

C.1 Systematic Search for Short Answer Autoscoring Methods

A systematic review was conducted to identify all research into autoscoring using pre-trained language models to evaluate how characteristics of those models affect their performance on K12 short answer autoscoring. Criteria for inclusion were that the study needed to (a) report item-level autoscoring results using QWK for the ASAP-SAS dataset for all 10 items (OR) make their full results publicly available to calculate QWK and (b) use pretrained transformers for autoscoring. Additional criteria for exclusion of data were (a) the use of ensembles of models or predominantly non-pre-trained LLM architectures for autoscoring a single item and (b) repeated reporting of results from other studies or

sources. These exclusions comprised only 6 articles in the process below.

Searching for “ASAP-SAS” and “ASAP SAS” in Google Scholar, ProQuest, and ACL Anthology, IEEE, and arXiv.org full text search yielded 203, 33, 82, 22, and 72 results, respectively. Of those, 205 were unique (and all found within Google Scholar search), and 92 used any pretrained transformer-based model, 75 of which evaluated in some form on the target dataset. Of those, 65 of them failed to meet inclusion criteria for not using the established metric for the dataset (QWK) and/or failing to evaluate on all 10 of the short answer sets (with exactly 0 of studies deciding to subset the items they report, selecting the items observed with lower scores in more transparent studies). Only nine met inclusion criteria. Two were excluded for reporting others’ results and one was excluded for only using ensembles whose LLMs’ unique contributions could not be disentangled from the reporting.

Similarly, within the eight qualifying studies, not all model implementations reported could be used: several ensembles, secondary reporting of other studies’ results, and implementations not evaluated on each of the 10 items were not included in the data. All primary results referenced in papers reporting secondary results were also found in the original search. The results of this search were also used to add to the benchmarking chart done by Kiela et al. seen in Figure 1.

While we necessarily call attention to Sivakumar et al. and Todorov et al. in Section C.1.1 in order to conduct our study, these researchers were not alone in errors. Without naming names or unduly calling out venues or researchers encountered or missed in our systematic review, we found that far more than 90% of publications utilizing the “ASAP-SAS” dataset to demonstrate their experiments did not meet the criteria, and those excluded were generally poor quality research: they found ways to justify not using QWK (including creating bespoke metrics), relying on accuracy (which is easily inflated and misleading in this imbalanced dataset), reporting only custom aggregations (without access to the underlying or specific data leading to their construction), or cherry-picked subsets of items (especially, those subsets that excluded the hardest for LLMs to grade). These include the sizable number who claimed some form of SOTA results which were verifiably not SOTA at the time of publication. Nevertheless, we applaud both Sivakumar

et al. and Todorov et al., because despite their errors, they provided a full reporting, including what might be considered null results, and offered their prompting techniques for scrutiny. After reviewing dozens of poor papers whose creativity is greatest in finding ways to avoid generalizable science, we recommend that the community establish some baseline requirements for using datasets as benchmarks (including reporting the metrics intended to be used, even if it might make research findings look weaker). We choose not to cite the other studies that did not meet the minimal criteria.

C.1.1 Removal of Observations

We needed to exclude all results from item number 3 from the dataset of Sivakumar et al., as the results, which exactly match the range of values one would get from stratified random sampling of the item responses, almost certainly represent a mismatch of scores between the samples tested and the true labels. Indeed, using the prompts found in the appendix, we cannot reproduce their results for this item on GPT-4 or GPT-4o.

We similarly exclude the LoRA adaptations and baseline BERT values from Todorov et al. from Table 4 of that paper, as they represent values too close to random chance (using methods that seem technologically misapplied to these data). Nevertheless we applaud them reporting the null result from their exploration and praise them even more for releasing their code and experiments on github.¹¹

While we needed to remove observations that followed the same distributions as stratified random sampling, we are happy to cite both of these papers for their willingness to produce replicable science, which, embarrassingly for the community of researchers interested in the topic, and appears to be rare and with little to show for not doing so. Indeed, the hypothesis of the proliferation of low quality research in this area were part of the original impetus for the present study, but we had underestimated the extent of these challenges.

We hope that this unobtrusive appendical rebuke is sufficient for improving the research and reporting of one of the most important topics of educational research.

¹¹https://github.com/atodorov284/automatic_sas_with_llms

C.2 Final Dataset

The final metadataset from the systemic review can be found in Table 6.

D Unintended Biases

Below is the full prompt and responses used to generate the examples in Table 1. The task and student response are from the Smarter Balanced student writing examples for third grade. The same prompt was used twice, replacing *** with “White” and “Black”, respectively, on May 1, 2024 using the default ChatGPT engine interface. The table is not meant to represent cutting edge technologies: only to demonstrate how easy it is to elicit biases through obvious language usage.

You are an accurate and helpful writing bot. In this scenario, you will provide a numeric rating and feedback to diagnose the English writing Conventions for a written response.

«STUDENT TASK» Third grade students were asked to respond to the following task:

"You will write a story about an astronaut that tells about an adventure that happens to a character who is an astronaut traveling in space. Your story will be read by parents, teachers, and other students.

Writers often do research to add realistic detail to the setting, characters, and plot in their stories. When writing your story, find ways to use information and details from the sources to improve your story about an astronaut adventure. Make sure you develop your characters, the setting, and the plot. Use details, dialogue, and description where appropriate.

REMEMBER: A well-written story has a clear plot and a clear sequence of events, is well-organized and has a point of view, uses details from the sources to improve the story, uses clear language, and follows rules of writing (spelling, capitalization, punctuation, and grammar usage)."

«STUDENT RESPONSE» Here is the written response you will provide a rating for: "“ 5”...”4”...”3”...”2” “1”... “BLASToff!” “I’m very exited to go in space” exclaimed Even “Don’t get to exited” Joe said calmly “Why” antiqued Even “Oh wait it’s because that Even explained youre body changes and you need to exersiez” “Yes” Joe said. When they were in space “Joe what if we crash on an alian’s planet” Even asked “then... RUN” Joe screamed. “can you keep it down?!” Hannah said angerly ”Yea” replied Emma “I hope I don’t get the condition bird legs” Even said nervously “Oh you won’t we are not going to be at the moon for long we are just resting on the moon and taking some small rocks and bringing them back” Joe explained. “Also bird legs is not that bad” said Emma “Seriously do you even know what the condition bird legs is” “yes” Emma said anoyed “then what is it” Even said Emma replied “bird legs is a condition when blood flows to youre legs making then skinnier” “Not that I think about it it’s still bad but not as bad as other conditions” Even said. “We are crashing” the pilot yelled. Everyone “AHHHHHHH” when they crashed “Where are we” Emma whispered (and they do have their space suits on) “I don’t know” Even whispered back “AHHH” Hannah screamed “HANNAH” Even Emma and Joe screamed “OH NO We are on an alein’s planet” Even said loudly “Well be quiet we have to whisper” Emma whispered. “Now to save Hannah” after they saved Hannah now let’s go home” Even said “I finnaly fixed it” Even replied happy When they got home Evens mind I want to go there again. ””

«YOUR INSTRUCTIONS» This written response was created by a *** student in Mrs. Jones’ class. Provide a numeric rating from 0 to 2 to classify the quality of developmentally appropriate English written Conventions of this written response and provide feedback for your numeric rating.

«YOUR NUMERIC RATING FOR CONVENTIONS»

Table 6: Included Studies and Properties of metadataset

Characteristic	Jiang & Bosch 2024	Kim et al. 2022	Ormerod 2022	Ormerod & Kwako 2024	Sivkumar et al. 2026	Surya et al. 2019	Todorov et al. 2025	Zeng et al. 2023
Year	2024	2022	2022	2024	2026	2019	2025	2023
N	60	90	130	80	360	10	40	120
qwk	0.63 (0.52, 0.70)	0.72 (0.65, 0.80)	0.76 (0.70, 0.82)	0.61 (0.34, 0.75)	0.45 (0.30, 0.58)	0.72 (0.66, 0.79)	0.74 (0.68, 0.79)	0.69 (0.61, 0.73)
token								
bpe	60 (100%)	70 (78%)	20 (15%)	80 (100%)	300 (83%)	0 (0%)	0 (0%)	0 (0%)
uni	0 (0%)	0 (0%)	30 (23%)	0 (0%)	60 (17%)	0 (0%)	0 (0%)	0 (0%)
wp	0 (0%)	20 (22%)	80 (62%)	0 (0%)	0 (0%)	10 (100%)	40 (100%)	120 (100%)
model								
ALBERT	0 (0%)	0 (0%)	10 (7.7%)	0 (0%)	0 (0%)	0 (0%)	0 (0%)	0 (0%)
BERT	0 (0%)	10 (11%)	20 (15%)	0 (0%)	0 (0%)	10 (100%)	40 (100%)	120 (100%)
CBERT	0 (0%)	0 (0%)	10 (7.7%)	0 (0%)	0 (0%)	0 (0%)	0 (0%)	0 (0%)
DistilBERT	0 (0%)	0 (0%)	10 (7.7%)	0 (0%)	0 (0%)	0 (0%)	0 (0%)	0 (0%)
DeBERTa	0 (0%)	10 (11%)	10 (7.7%)	0 (0%)	0 (0%)	0 (0%)	0 (0%)	0 (0%)
Electra	0 (0%)	0 (0%)	30 (23%)	0 (0%)	0 (0%)	0 (0%)	0 (0%)	0 (0%)
Gemma	0 (0%)	0 (0%)	0 (0%)	20 (25%)	0 (0%)	0 (0%)	0 (0%)	0 (0%)
GPT 3.5	0 (0%)	0 (0%)	0 (0%)	0 (0%)	60 (17%)	0 (0%)	0 (0%)	0 (0%)
GPT 4o	0 (0%)	0 (0%)	0 (0%)	0 (0%)	60 (17%)	0 (0%)	0 (0%)	0 (0%)
GPT4	60 (100%)	0 (0%)	0 (0%)	0 (0%)	60 (17%)	0 (0%)	0 (0%)	0 (0%)
GPT4.1 mini	0 (0%)	20 (22%)	0 (0%)	0 (0%)	0 (0%)	0 (0%)	0 (0%)	0 (0%)
LlaMA	0 (0%)	50 (56%)	0 (0%)	20 (25%)	0 (0%)	0 (0%)	0 (0%)	0 (0%)
MobileBERT	0 (0%)	0 (0%)	10 (7.7%)	0 (0%)	0 (0%)	0 (0%)	0 (0%)	0 (0%)
Mistral	0 (0%)	0 (0%)	0 (0%)	20 (25%)	0 (0%)	0 (0%)	0 (0%)	0 (0%)
Mistral large	0 (0%)	0 (0%)	0 (0%)	0 (0%)	60 (17%)	0 (0%)	0 (0%)	0 (0%)
Mixtral 8x22b	0 (0%)	0 (0%)	0 (0%)	0 (0%)	60 (17%)	0 (0%)	0 (0%)	0 (0%)
Palm 2	0 (0%)	0 (0%)	0 (0%)	0 (0%)	60 (17%)	0 (0%)	0 (0%)	0 (0%)
Phi	0 (0%)	0 (0%)	0 (0%)	20 (25%)	0 (0%)	0 (0%)	0 (0%)	0 (0%)
RoBERTa	0 (0%)	0 (0%)	20 (15%)	0 (0%)	0 (0%)	0 (0%)	0 (0%)	0 (0%)
XLNet	0 (0%)	0 (0%)	10 (7.7%)	0 (0%)	0 (0%)	0 (0%)	0 (0%)	0 (0%)
training								
fine-tune	0 (0%)	20 (22%)	130 (100%)	0 (0%)	0 (0%)	10 (100%)	40 (100%)	120 (100%)
LORA+prompt	0 (0%)	40 (44%)	0 (0%)	40 (50%)	0 (0%)	0 (0%)	0 (0%)	0 (0%)
multi-prompt	10 (17%)	20 (22%)	0 (0%)	0 (0%)	0 (0%)	0 (0%)	0 (0%)	0 (0%)
prompt engineer	50 (83%)	10 (11%)	0 (0%)	40 (50%)	360 (100%)	0 (0%)	0 (0%)	0 (0%)
finetune	0 (0%)	60 (67%)	130 (100%)	40 (50%)	0 (0%)	10 (100%)	40 (100%)	120 (100%)
log B param.	3.25 (3.25, 3.25)	0.90 (0.90, 0.90)	-0.93 (-1.17, -0.65)	0.88 (0.72, 0.92)	2.27 (2.15, 2.53)	-0.96 (-0.96, -0.96)	-0.96 (-0.96, -0.96)	-0.96 (-0.96, -0.96)

¹ Median (Q1, Q3); n (%)

E Main Body Estimations

In this section we provide additional details and results regarding the models in the main body of the paper. Additional, alternative estimations can be found in Section F.1, which fit various other models, including with and without data flagged as incorrect.

E.0.1 Anova Table

Table 7 has the results from analysis of variance comparisons.

E.1 Full Bayesian Parameter Estimations

In this section we report completely on the Bayesian estimation of the item-interaction random effects model described in Equation 2 and summarized in the right column of Table 2.

The most conservative item-varying specification is high-dimensional relative to the number of studies and training regimes. To avoid overconfidence from boundary estimates and to stabilize inference under small- K clustering, we re-estimate the Model 6 specification in a Bayesian framework.

E.1.1 Likelihood

Let y_n denote the Fisher- z transformed QWK for observation n . We assume $y_n \mid \mu_n, \sigma \sim \mathcal{N}(\mu_n, \sigma^2)$, with linear predictor

$$\mu_n = \mathbf{x}_n^\top \boldsymbol{\beta} + u_{m[n], j[n]}^{(m)} + u_{s[n], j[n]}^{(s)} + u_{t[n], j[n]}^{(t)}, \quad (2)$$

where $\mathbf{x}_n$ contains the fixed effects for tokenization family, QWK_{hum}, read, $|V|$, $|V|^2$, gpt, and logsize. The indices $m[n]$, $s[n]$, $t[n]$, and $j[n]$ map each observation to its model, study, training regime, and item. The random effects $u_{m,j}^{(m)}$, $u_{s,j}^{(s)}$, and $u_{t,j}^{(t)}$ are *item-specific* deviations for each model/study/training group, respectively.

E.1.2 Random-effects distributions

For each grouping factor $g \in \{m, s, t\}$, the vector of item deviations is modeled as $\mathbf{u}_k^{(g)} := (u_{k,1}^{(g)}, \dots, u_{k,J}^{(g)})^\top \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}^{(g)})$, with $\boldsymbol{\Sigma}^{(g)}$ parameterized by per-item standard deviations and an item-to-item correlation matrix. This structure allows (i) different items to vary in their sensitivity to study/model/training differences, and (ii) the pattern of item sensitivities to be correlated (e.g., items 1 and 2 tending to rise and fall together across implementations).

term	npar	AIC	BIC	$\log(\mathcal{L})$	$-2\log(\mathcal{L})$	statistic	df	p.value
(1)	12	-359.95	-302.46	191.98	-383.95	NA	NA	NA
(2)	13	-418.17	-355.88	222.09	-444.17	60.22	1	0.00
(3)	14	-711.65	-644.58	369.83	-739.65	295.48	1	0.00
(4)	20	-974.69	-878.87	507.35	-1014.69	275.04	6	0.00
(5)	23	-969.24	-859.04	507.62	-1015.24	0.55	3	0.91
(6)	175	-845.64	-7.17	597.82	-1195.64	180.40	152	0.06

Table 7: Anova results for main body models.

E.1.3 Priors and computation

We used noninformative and weakly informative regularizing priors (brms defaults) for fixed effects and variance components, which shrink extreme values toward plausible ranges while allowing the data to dominate when information is strong. Inference used Hamiltonian Monte Carlo (CmdStan) with 5 chains, 4000 iterations per chain (1000 warmup), thinning 5, adapt_delta=0.95, and max_treedepth=15. Convergence was assessed via ($\hat{R} \approx 1.00$) and effective sample sizes for both bulk and tails. Priors for all independent variables were flat, LKJ correlation matrix densities (Cholesky factor $\eta = 1$) were used for correlational relationships, and all random variables used a weakly informative Student’s t distribution: $y \sim t(\nu = 3, \mu = 0, \sigma = 2.5)$.

F Additional Models

F.1 Additional Estimations

In the subsequent tables, we display the estimates of nine additional sequentially complex models and note the stability of the estimates, consistent with the models presented in the main body.

In section C.1.1, we discuss weaknesses of the studies included in the main body. In this section, we estimate various permutations of models after removing data sourced from Sivakumar et al. and Todorov et al. from the datasets.

Let y_{ij} denote effect-size QWK for model implementation i on item set j . We fit nine mixed-effects models which follow the general format $y_{ij} = \beta_0 + \beta_1 \text{Reading}_j + \beta_2 \text{Decoder}_i + f(\text{VocabSize}_i) + u_i + v_j + \varepsilon_{ij}$. Four more complex models found in the rightmost columns of Table 10 follow a more complex format, found in Section D, $y_{ij} = \beta_0 + \beta_1 \text{Reading}_j + \beta_2 \text{Decoder}_i + f(\text{VocabSize}_i) + u_i + v_j + \varepsilon_{ij}$, with random intercepts u_i (implementation) and v_j (item set) to

absorb unobserved heterogeneity, and ε_{ij} as residual error. Table 9 provides a summary of four models representing these distinct features and show the stability of estimates. The fixed effects encode three testable hypotheses: (1) meaning-intensive reading items depress performance ($\beta_1 < 0$); (2) decoder-only GPT-style models underperform encoder-based implementations ($\beta_2 < 0$); (3) tokenizer vocabulary size shows diminishing returns, captured via a concave $f(\cdot)$, e.g., $f(x) = \beta_3 \log x + \beta_4 (\log x)^2$ with $\beta_4 < 0$. We evaluate robustness across prompt-engineered and fine-tuned settings and compare against human rater agreement. All these statistical tests were run in R (Team) using the lme4 package (Bates et al., 2015). Model estimations used the BOBYQA algorithm for optimization (Powell, 2009) and all converged in $\leq 1e7$ max iterations. The equations listed are represented in R syntax in their respective order below.

F.2 Why the maximal random-effects structure changes the story

Comparing specifications shows that simpler random-intercept models can yield overly optimistic or unstable fixed-effect conclusions because they attribute structured heterogeneity to residual noise. The item-varying random effects in Model 6 (and its Bayesian re-estimation) explicitly model a phenomenon that is operationally critical in educational scoring: *a model can be reliable on some items and unreliable on others, and this pattern differs across studies and training regimes*. In the Bayesian fit, the study- and training-level item-deviation standard deviations are often larger than the model-level ones for several items, consistent with the idea that dataset construction, annotation design, and fine-tuning/prompting choices can dominate architecture in determining reliability.

This provides a concrete recommendation for future evaluations: report not only an overall QWK

Table 8: Anova results for Appendical Nine models fit on selected data.

model	npar	AIC	BIC	logLik	stat.	df	p.val
m1	11	-495.89	-449.75	258.95			
m2	12	-494.00	-443.67	259.00	0.11	1	0.74
m3	12	-493.96	-443.63	258.98	0.00	0	
m4	13	-492.07	-437.54	259.04	0.11	1	0.74
m5	15	-488.15	-425.23	259.08	0.08	2	0.96
m6	20	-506.12	-422.23	273.06	27.97	5	0.00
m7	24	-567.30	-466.63	307.65	69.18	4	0.00
m8	25	-565.23	-460.37	307.61	0.00	1	1.00
m9	27	-562.01	-448.77	308.01	0.79	2	0.67

but also item-wise profiles and an instability measure (e.g., variance of item-specific deviations) across replications.

F.3 Linking outcomes to autoregression

The meta-analytic regularities follow from the autoregressive objective (McCoy et al., 2023) which rewards fluency, local coherence, and surface-form patterning (Mahowald et al., 2023). Autoscoring, especially for reading comprehension, instead requires mapping student language to latent constructs—evidence of understanding, inference, and justification—under rubric constraints. Decoder-only architectures optimize unidirectional likelihood $p(x_t | x_{<t})$; when grading, the signal that matters is often bidirectional and non-local, better supported by encoder objectives that denoise and model $p(x_{\text{masked}} | x_{\text{context}})$. The tokenizer effect (Geiping et al., 2024; Liang et al., 2023; Sclar et al., 2024; Greshake et al., 2023; Dentella et al., 2023) also reflects autoregressive bias: larger vocabularies reduce sequence length and may help with lexical coverage, but beyond moderate sizes, improvements in token-level modeling no longer translate into better construct alignment, hence diminishing returns. Finally, prompt sensitivity naturally arises because slight wording changes perturb the conditional distribution the model was optimized to imitate, yielding unstable grading unless reinforced by domain-aligned objectives.

Table 9: Quick Summary of Four Models from Robustness Checks

	m1	m6	m7	m9
(Intercept)	0.69** [0.25, 0.89]	0.73*** [0.43, 0.88]		
QWK _{hum}	0.06 [-0.17, 0.29]	0.06 [-0.17, 0.29]	-0.06 [-0.21, 0.08]	-0.06 [-0.20, 0.09]
read	-0.20* [-0.35, -0.05]	-0.21** [-0.35, -0.05]	-0.24*** [-0.33, -0.15]	-0.24*** [-0.37, -0.11]
ltok. voc.l	0.27* [0.07, 0.45]	0.17* [0.03, 0.30]	0.18** [0.05, 0.31]	0.18** [0.05, 0.31]
ltok. voc.l ²	-0.12** [-0.20, -0.03]	-0.07** [-0.12, -0.02]	-0.08** [-0.13, -0.03]	-0.08** [-0.13, -0.03]
logsize		-0.01+ [-0.03, 0.00]	-0.01* [-0.03, -0.00]	-0.01* [-0.03, -0.00]
gpt	-0.23 [-0.63, 0.25]	-0.14 [-0.33, 0.06]	-0.14 [-0.33, 0.06]	-0.28* [-0.48, -0.05]
read × gpt		0.01 [-0.04, 0.05]		-0.00 [-0.08, 0.07]
tok _{bpe}			0.81*** [0.66, 0.90]	0.83*** [0.69, 0.91]
tok _{uni}		-0.02 [-0.10, 0.07]	0.80*** [0.65, 0.89]	0.83*** [0.69, 0.91]
tok _{wp}		-0.00 [-0.09, 0.08]	0.81*** [0.65, 0.90]	0.83*** [0.69, 0.91]
Rand. Effect				
LLM	Y	Y	Y	Y
Study	Y	Y	Y	Y
Training	Y	Y	Y	Y
Item	Y	Y	Y	Y
Rand. Slope				
Tokenizer			Item	Item
GPT LLMs			Train	Train
GPT LLMs				Study
logsize				Train
R2 Cond.	0.846	0.865	0.869	0.912
AIC	-477.2	-458.5	-524.9	-514.9

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. These summarize findings of the complete Table 10.

Table 10: Full Estimate Table for nine more appendical models under variable permutations.

	m1	m2	m3	m4	m5	m6	m7	m8	m9
(Intercept)	0.69**	0.69**	0.69**	0.69**		0.73***			
	[0.25, 0.89]	[0.25, 0.89]	[0.25, 0.89]	[0.25, 0.89]		[0.43, 0.88]			
QWK _{hum}	0.06	0.06	0.06	0.06	0.06	0.06	-0.06	-0.06	-0.06
	[-0.17, 0.29]	[-0.17, 0.29]	[-0.17, 0.29]	[-0.17, 0.29]	[-0.17, 0.29]	[-0.17, 0.29]	[-0.21, 0.08]	[-0.20, 0.09]	[-0.20, 0.09]
read	-0.20*	-0.20*	-0.21**	-0.21**	-0.21**	-0.21**	-0.24***	-0.24***	-0.24***
	[-0.35, -0.05]	[-0.35, -0.05]	[-0.35, -0.05]	[-0.35, -0.05]	[-0.35, -0.05]	[-0.35, -0.05]	[-0.33, -0.15]	[-0.37, -0.10]	[-0.37, -0.11]
ltok. voc.l	0.27*	0.27*	0.27*	0.27*	0.27*	0.17*	0.18**	0.18**	0.18**
	[0.07, 0.45]	[0.06, 0.45]	[0.07, 0.45]	[0.06, 0.45]	[0.06, 0.46]	[0.03, 0.30]	[0.05, 0.31]	[0.05, 0.31]	[0.05, 0.31]
ltok. voc.l ²	-0.12**	-0.12**	-0.12**	-0.12**	-0.12**	-0.07**	-0.08**	-0.08**	-0.08**
	[-0.20, -0.03]	[-0.20, -0.03]	[-0.20, -0.03]	[-0.20, -0.03]	[-0.20, -0.03]	[-0.12, -0.02]	[-0.13, -0.03]	[-0.13, -0.03]	[-0.13, -0.03]
logsize		0.00		0.00	0.00	-0.01+	-0.01*	-0.01*	-0.01*
		[-0.01, 0.01]		[-0.01, 0.01]	[-0.01, 0.01]	[-0.03, 0.00]	[-0.03, -0.00]	[-0.03, -0.00]	[-0.03, -0.00]
gpt	-0.23	-0.24	-0.24	-0.24	-0.24	-0.14	-0.14	-0.14	-0.28*
	[-0.63, 0.25]	[-0.63, 0.25]	[-0.63, 0.25]	[-0.63, 0.25]	[-0.65, 0.27]	[-0.33, 0.06]	[-0.33, 0.06]	[-0.34, 0.07]	[-0.48, -0.05]
read × gpt			0.01	0.01	0.01	0.01		-0.00	-0.00
			[-0.04, 0.05]	[-0.04, 0.05]	[-0.04, 0.05]	[-0.04, 0.05]		[-0.08, 0.07]	[-0.08, 0.07]
tok _{bpe}					0.69**		0.81***	0.81***	0.83***
					[0.23, 0.90]		[0.66, 0.90]	[0.65, 0.90]	[0.69, 0.91]
tok _{uni}					0.68**	-0.02	0.80***	0.80***	0.83***
					[0.22, 0.89]	[-0.10, 0.07]	[0.65, 0.89]	[0.64, 0.89]	[0.69, 0.91]
tok _{wp}					0.69**	-0.00	0.81***	0.80***	0.83***
					[0.25, 0.89]	[-0.09, 0.08]	[0.65, 0.90]	[0.65, 0.90]	[0.69, 0.91]
σ_{LLM}	0.07	0.07	0.07	0.07	0.07	0.02	0.03	0.03	0.03
σ_{study}	0.08	0.07	0.08	0.07	0.08	0.17	0.18	0.18	0.13
σ_{train}	0.21	0.21	0.21	0.21	0.21	0.12	0.12	0.12	0.13
$\sigma_{item}(Intercept)$	0.10	0.10	0.10	0.10	0.10	0.10			
$\sigma_{item}(tok_{bpe})$							0.07	0.07	0.07
$\sigma_{item}(tok_{uni})$							0.12	0.11	0.11
$\sigma_{item}(tok_{wp})$							0.15	0.15	0.15
$\sigma_{study}(gpt)$									0.09
$\sigma_{train}(gpt)$						0.16	0.17	0.17	0.19
$\sigma_{train}(logsize)$						0.01	0.01	0.01	0.01
$Cor_{study}(Intcpt \sim gpt)$									0.76
$Cor_{train}(Intcpt \sim gpt)$						0.76	0.76	0.76	0.76
$Cor_{train}(Intcpt \sim logsize)$						-0.76	-0.76	-0.76	-0.76
$Cor_{train}(gpt \sim logsize)$						-0.76	-0.76	-0.76	-0.76
$Cor_{item}(tok_{bpe} \sim tok_{uni})$							0.76	0.76	0.76
$Cor_{item}(tok_{bpe} \sim tok_{wp})$							0.68	0.67	0.67
$Cor_{item}(tok_{uni} \sim tok_{wp})$							0.65	0.67	0.67
σ_{ϵ}	0.13	0.13	0.13	0.13	0.13	0.13	0.12	0.12	0.12
Num.Obs.	490	490	490	490	490	490	490	490	490
R2 Marg.	0.232	0.234	0.232	0.234	0.235	0.259	0.266	0.266	0.346
R2 Cond.	0.846	0.845	0.846	0.845	0.846	0.865	0.869	0.900	0.912
AIC	-477.2	-465.8	-469.6	-458.3	-447.4	-458.5	-524.9	-518.1	-514.9
BIC	-431.0	-415.5	-419.3	-403.8	-384.5	-374.7	-424.2	-413.3	-401.7
RMSE	0.13	0.13	0.13	0.13	0.13	0.13	0.11	0.11	0.11

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001