

EMNLP 2026

**The 2026 Conference on Empirical Methods in Natural
Language Processing (EMNLP 2026)**

Tutorial Abstracts

October 24-29, 2026

©2026 Association for Computational Linguistics

Order copies of this and other ACL proceedings from:

Association for Computational Linguistics (ACL)
317 Sidney Baker St. S
Suite 400 - 134
Kerrville, TX 78028
USA
Tel: +1-855-225-1962
acl@aclweb.org

ISBN 979-8-89176-463-7

Introduction

Welcome to the EMNLP 2026 tutorials!

We believe that our tutorial program reflects where NLP research has been heading in 2026 and will continue in 2027. Some sessions dig deep into the technical core of the field, for example, how we pre-process and feed data into models, design agents, and push models toward formalization, while others step back to consider bigger questions, like how LLM-based technology compares to humans and what LLM-driven progress means for global inclusion. No set of few tutorials could capture everything our community cares about, but we think this selection gives a good sense of the conversations you'll encounter throughout the conference and beyond.

This year, we did a joint call for tutorial proposals with ACL-IJCNLP alongside their tutorial co-chairs: Naoaki Okazaki and Zhongqing Wang. We received 57 submissions of which 8 were accepted to ACL and 6 are taking place here at EMNLP. All of the proposals were reviewed by multiple of the tutorial co-chairs before a broader discussion among all. As in previous years, it was a highly competitive selection process with many excellent submissions. We fully acknowledge that 6 tutorials (the number was a hard limit due to venue constraints) by no means capture all of the potentially interesting topics relevant to our research community, but we hope that this selection will engage attendees from a wide range of backgrounds and expertise levels after the main conference closes.

We highly appreciate the efforts of all the tutorial authors who provide compelling lectures and talks and make this day possible. We would also like to extend a special thanks to Jenn Rachford who works tirelessly behind the scenes to help organize all parts of the conference, including the tutorials.

Enjoy the conference and, especially, the tutorials!

Ana Marasović and Goran Glavaš
EMNLP 2026 Tutorial Co-Chairs

Organizing Committee

General Chair

André Martins, Instituto Superior Técnico, University of Lisbon

Program Chairs

Hung-Yi Lee, National Taiwan University

Sunipa Dev, Google

Jessy Li, University of Texas at Austin

Desmond Elliott, University of Copenhagen

Tutorial Chairs

Ana Marasović, University of Utah

Goran Glavaš, University of Würzburg

Table of Contents

<i>Failure Modes in Agentic AI: From Reproducible Triggers and Diagnostics to Reliable Fixes</i> Zoey Sha Li, Zihan Wang, Hejie Cui and Manling Li	1
<i>Language Models as Windows into Human Language Processing: Neural Encoding, Decoding, and Brain Alignment with NLP</i> Jingyuan Sun, Yunhao Zhang, Yifan Wang and Yang Cui	2
<i>Tokenization in Large Language Models</i> Valentin Hofmann, Julie Kallini, Sachin Kumar, Jindřich Libovický, Tomasz Limisiewicz and Elizabeth Salesky	4
<i>Autoformalization & Mathematical Reasoning</i> Lan Zhang, Marco Valentino, Jordan Meadows and André Freitas	6
<i>Causal Foundations for Interpretability: What Your Methods Actually Identify</i> Shruti Joshi, Aaron Mueller, Patrik Reizinger and Dhanya Sridhar	8
<i>Multilingual and Multicultural LLMs</i> Wei Xu, David Ifeoluwa Adelani, Vinodkumar Prabhakaran, Sara Hooker and Alan Ritter . . .	10

Program

Wednesday, October 28, 2026

- | | |
|---------------|---|
| 09:00 - 12:30 | <i>Failure Modes in Agentic AI: From Reproducible Triggers and Diagnostics to Reliable Fixes</i> |
| 09:00 - 12:30 | <i>Language Models as Windows into Human Language Processing: Neural Encoding, Decoding, and Brain Alignment with NLP</i> |
| 09:00 - 12:30 | <i>Tokenization in Large Language Models</i> |
| 14:00 - 17:30 | <i>Autoformalization & Mathematical Reasoning</i> |
| 14:00 - 17:30 | <i>Causal Foundations for Interpretability: What Your Methods Actually Identify</i> |
| 14:00 - 17:30 | <i>Multilingual and Multicultural LLMs</i> |

Failure Modes in Agentic AI: From Reproducible Triggers and Diagnostics to Reliable Fixes

Zoey Sha Li, Zihan Wang, Hejie Cui, Manling Li

<https://failure-mode-agent.github.io/>

Foundation-model agents are being deployed faster than the community can characterize how they fail. In barely two years they have moved into browsing, software engineering, computer use, scientific discovery, and embodied control, where reliability is governed by long-horizon, closed-loop interaction that current single-turn evaluation paradigms rarely capture. Failures (error cascades, reward hacking, mode and reasoning collapse, memory drift, and over-reasoning loops) still lack a shared vocabulary, reproducible triggers, and comparable diagnostics. This cutting-edge tutorial gives a systematic, mechanism-first account of agentic failure modes across four axes: (1) environments and infrastructure, (2) training dynamics, (3) evaluation and red teaming, and (4) test-time scaling and adaptation, with safety threaded throughout. Attendees leave with reproducible triggers, comparable diagnostic protocols, and verifiable mitigations.

Zoey Sha Li

email: mail@zoeyli.com

website: zoeyli.com

Zoey Li is a Senior Applied Scientist at Amazon. She works on pretraining scaling, data selection, and post-training diagnostics. She received her Ph.D. from the University of Illinois Urbana-Champaign in 2024, advised by Jiawei Han and Heng Ji.

Zihan Wang

email: zihanw@u.northwestern.edu

website: <https://zihanwang314.github.io/>

Zihan Wang is a PhD student at Northwestern University, advised by Manling Li. He works on reinforcement learning for reasoning agents in multi-turn, stochastic environments. He is the lead author of the RAGEN series, which characterizes core agentic failure modes.

Hejie Cui

email: hejie.cui@stanford.edu

website: <https://hejiecui.com>

Hejie Cui is a Senior Research Scientist at Google. She works on large language models and agentic AI systems. She received her Ph.D. in Computer Science from Emory University, advised by Carl Yang, and was a postdoctoral researcher at Stanford University with Nigam H. Shah and Sanmi Koyejo, focused on LLM post-training and evaluation.

Manling Li

email: manling.li@northwestern.edu

website: <https://limanling.github.io/>

Manling Li is an Assistant Professor at Northwestern University, and Amazon Scholar. Her research focuses on grounded reasoning and multimodal knowledge for foundation-model agents, recognized by the ACL 2025 Dissertation Award Honorable Mention and paper awards from ACL and NAACL. She co-leads multiple recent open-source efforts on agentic AI, and has organized community events around foundation models and multimodal reasoning.

Language Models as Windows into Human Language Processing: Neural Encoding, Decoding, and Brain Alignment with NLP

Jingyuan Sun, Yunhao Zhang, Yifan Wang, Yang Cui

<https://www.ai4bci.com/affiliated-events/emnlp2026tutorial>

Modern language models are not only tools for processing text; they are also becoming computational instruments for studying how humans understand, predict, and generate language. This tutorial introduces neural encoding and neural decoding to the NLP community through this broader perspective: how can NLP help us understand language processing in the human brain, and how can brain data help us evaluate and improve language models? We will present neural encoding as a way to test whether representations from language models—such as contextual embeddings, surprisal, attention patterns, and layer-wise features—can explain brain responses during reading or listening. We will then introduce neural decoding as the reverse problem: reconstructing linguistic or multimodal content from brain activity using retrieval, representation learning, and foundation-model-guided generation. Rather than framing decoding mainly as brain–computer interface technology, we will connect it to core NLP questions about meaning representation, generation under uncertainty, multimodal grounding, and semantic evaluation. The tutorial assumes no neuroscience background. It combines intuitive foundations, representative papers, practical modeling recipes, notebook-style demos, and discussion of mental privacy, consent, and responsible brain–language technologies.

Jingyuan Sun

email: jingyuan.sun@manchester.ac.uk

website: https://sites.google.com/view/jingyuan-sun/home?read_current=1

Dr. Jingyuan Sun is an Assistant Professor in the Department of Computer Science at the University of Manchester. His research lies at the intersection of artificial intelligence, computational linguistics, and cognitive neuroscience. He develops multimodal models for neural encoding and decoding, including sentence-level brain encoding, fMRI-to-image reconstruction, and semantic video decoding. His work has appeared in venues including NeurIPS, AACL, NAACL, ACL tutorial tracks, and IEEE TNNLS. He has delivered international tutorials at ACL 2024 and IJCAI 2024 and has experience teaching interdisciplinary audiences in NLP, machine learning, and neuroscience.

Yunhao Zhang

email: zhangyunhao2021@ia.ac.cn

Dr. Yunhao Zhang is a Research Assistant Professor at the Institute of Automation, Chinese Academy of Sciences. His research lies at the intersection of artificial intelligence, cognitive neuroscience, and psycholinguistics. He uses large language models to analyze multimodal neuroimaging data, with the goal of investigating the neural mechanisms underlying semantic representation in the human brain. His related work has appeared in venues including NeuroImage, Scientific Data, Language Resources and Evaluation, EMNLP, and CogSci.

Yifan Wang

email: yifan.wang@manchester.ac.uk

Yifan Wang is a PhD researcher at the University of Manchester with hands-on experience in multimodal learning, biomedical data processing, and neural model implementation. In this tutorial, he will support the practical demonstrations and Q&A, particularly for the neural decoding component.

Yang Cui

email: `yang.cui@manchester.ac.uk`

Yang Cui is a PhD researcher at the University of Manchester with hands-on experience in multimodal learning, biomedical data processing, and neural model implementation. In this tutorial, he will support the practical demonstrations and Q&A, particularly for the neural encoding and brain-alignment component.

Tokenization in Large Language Models

Valentin Hofmann, Julie Kallini, Sachin Kumar, Jindřich Libovický, Tomasz Limisiewicz,
Elizabeth Salesky

<https://tokenization-tutorial.github.io>

Tokenization defines how data is represented as input and output for machine learning systems, including large language models (LLMs). While tokenization has long been a foundational component of NLP systems, it has historically remained a niche research area, often considered solved. The recent proliferation of LLMs has changed this: as these models are deployed across diverse languages and domains, tokenization-related issues have become increasingly apparent—from cross-lingual representation disparities to efficiency bottlenecks. This tutorial will provide a comprehensive overview of tokenization methods in text, starting from classical word-based and morphological approaches, to the current standard of statistical subword tokenizers, and emerging methods on end-to-end learnable tokenizers. We will also discuss the implications of tokenization choices on model performance, fairness, and efficiency, and present best practices for evaluating and designing tokenizers in multilingual and domain-specific settings.

Valentin Hofmann

email: valentin.hofmann@lmu.de

website: <https://valentinhofmann.github.io>

Valentin Hofmann is an assistant professor at LMU Munich and a principal investigator at the Munich Center for Machine Learning. His research lies at the intersection of natural language processing, linguistics, and social science, with a particular focus on tokenization, evaluation, and fairness. He holds a Ph.D. in Linguistics from the University of Oxford and a Ph.D. in Computer Science from LMU Munich. Previously, he was a postdoctoral researcher at the Allen Institute for AI and the University of Washington.

Julie Kallini

email: kallini@stanford.edu

website: <https://juliekallini.com>

Julie Kallini is a Ph.D. student in the Computer Science Department at Stanford University, advised by Christopher Potts and Dan Jurafsky. Her research focuses on language model architectures and efficiency, computational linguistics, cognitive science, and tokenization. She has published first-author papers at ACL, EMNLP, ICLR, and ICML, and her work received the Best Paper Award at ACL 2024. She also has extensive teaching experience and served as the head teaching assistant for Stanford’s NLP course in 2026.

Sachin Kumar

email: kumar.1145@osu.edu

website: <https://www.shocheen.com/>

Sachin Kumar is an assistant professor at the Ohio State University. He received his Ph.D. at Carnegie Mellon University. His research focuses on building user-centric NLP models that work well for all languages, populations, and individuals, covering research areas such as tokenization unfairness, tokenizer adaptation, and dynamic/learnable tokenization. His work has been published at venues such as ACL, EMNLP, NAACL, EACL, ICLR, and NeurIPS. He has previously co-taught tutorials at WWW

2022 and EMNLP 2023, as well as co-organized workshops at EMNLP 2024, ICML 2025, COLM 2025, ACL 2026, and COLM 2026.

Jindřich Libovický

email: libovicky@ufal.mff.cuni.cz

website: <https://ufal.mff.cuni.cz/jindrich-libovicky>

Jindřich Libovický is a Research Associate at the Institute of Formal and Applied Linguistics (ÚFAL), Charles University, Prague, where he leads a research group on multilinguality and cross-lingual fairness. His research spans multilingual language modeling and machine translation, with particular focus on how language models align across languages, how to design better tokenization and subword segmentation methods for multilingual settings, and how model performance varies across languages in ways that affect fairness. He holds a Ph.D. in Computational Linguistics from Charles University and previously worked as a researcher at LMU Munich.

Tomasz Limisiewicz

email: tomlim@cs.washington.edu

website: <https://tomlimi.github.io>

Tomasz Limisiewicz is a Postdoctoral Researcher at Meta (FAIR) and the University of Washington. His research motivation is developing better and fairer (multilingual) language models. Tomasz's recent work explores more efficient data representation by language modeling, and latent tokenization, applicable across languages and further beyond text. His work has been published in ACL, EMNLP, NAACL, and ICLR conferences.

Elizabeth Salesky

email: elizabeth.salesky@gmail.com

website: <https://esalesky.github.io>

Elizabeth Salesky is a Research Scientist at Google DeepMind. She holds a Ph.D. from the Center for Language and Speech Processing at Johns Hopkins University, where her dissertation on robust open-vocabulary processing with pixel-based text representations won the 2026 AMTA Best Thesis Award. Her research focuses on multilingual speech and language processing, with a focus on robust cross-lingual and cross-modal transfer, visually-situated text, and speech translation. She has co-organized several workshops and a speech translation tutorial at EACL 2021.

Autoformalization & Mathematical Reasoning

Lan Zhang, Marco Valentino, Jordan Meadows, André Freitas

<https://sites.google.com/view/autoformalization-tutorial>

Recent advances in large language models, formal reasoning, theorem proving, and agentic AI are rapidly changing what can be augmented and automated in mathematics. Yet mathematical reasoning is considerably richer than producing correct final answers: it involves understanding mathematical discourse, constructing and revising arguments, moving between informal and formal representations, exploring alternative proof strategies, verifying intermediate claims, and ultimately supporting mathematical discovery. This tutorial provides a structured introduction to the emerging field of automated mathematical reasoning, with a particular focus on autoformalization and neuro-symbolic and agentic approaches. We first examine mathematical reasoning from linguistic and conceptual perspectives, decomposing the spectrum from informal mathematical discourse to fully formal proofs. We then introduce the principal computational mechanisms underlying modern systems, including formal representations, autoformalization, theorem provers, critics and proxy evaluators, refinement loops, and complex agentic architectures. The tutorial connects these foundations with practical system-building, evaluation methodologies, current state-of-the-art results, and demonstrations of existing frameworks. We further examine how mathematical reasoning methods can extend toward physics, where mathematical arguments must interact with physical models, assumptions, approximations, and scientific knowledge. Finally, we use ambitious mathematical and scientific problems to identify the capabilities still required to move from benchmark solving toward systems that can participate meaningfully in verification, reasoning, and discovery.

Lan Zhang

email: lan.zhang-6@postgrad.manchester.ac.uk

website: <https://lanzhang128.github.io/>

Lan is a PhD student in the Department of Computer Science in University of Manchester. His research interest mainly focuses on autoformalization and formal mathematical reasoning. He has published papers at major NLP conferences, including ACL, EMNLP, and NAACL. His contributions to autoformalization were recognized with the Best Resource Paper Award at EMNLP 2025, highlighting the impact of his work on formal mathematical reasoning and verification.

Marco Valentino

email: m.valentino@sheffield.ac.uk

website: <https://www.marcovalentino.net/>

Marco is a Lecturer in Artificial Intelligence and Applications of AI in the Natural Language Processing (NLP) group at the University of Sheffield. Prior to Sheffield, he was a member of the Neuro-Symbolic AI Group at the Idiap Research Institute in Switzerland, and obtained a PhD in Computer Science from the University of Manchester. His research focuses on developing AI systems that can use explanation as a core mechanism for learning and reasoning, investigating the integration of neural and symbolic AI methods. Moreover, he is interested in developing methodologies to interpret, control, and evaluate Large Language Models (LLMs), with a focus on disentangling knowledge acquisition from abstract logical reasoning, and enabling out-of-distribution, out-of-domain generalisation. His research on neuro-symbolic NLP received a best resource paper award at EMNLP 2025 and an outstanding paper award at EMNLP 2024.

Jordan Meadows

email: jmeadows@idiap.ch

Jordan Meadows is a Postdoctoral Researcher at Idiap Research Institute working on natural language processing, neuro-symbolic AI, and AI for Science. He received his PhD in Computer Science from the University of Manchester, where he studied mathematical language processing and the reasoning capabilities of language models. His research focuses broadly on mathematical and scientific reasoning, with particular interest in combining language models with symbolic methods to make AI systems more reliable, interpretable, and verifiable. His work has been published across major NLP and AI venues, and he has presented research on mathematical reasoning and the autoformalisation of science to both academic and interdisciplinary audiences.

André Freitas

email: andre.freitas@idiap.ch

website: <https://www.andrefreitas.net/>

André Freitas leads the Neuro-symbolic AI Lab at the University of Manchester and IDIAP Research Institute. His main research interests are on enabling the development of AI methods to support abstract, flexible and controlled reasoning in order to support AI-augmented scientific discovery. In particular, he investigates how the combination of neural and symbolic data representation paradigms can deliver better models of inference. He is an active contributor to the main conferences and journals in the AI/Natural Language Processing (NLP) interface (AAAI, NeurIPS, ACL, EMNLP, EACL, COLING, TACL, Computational Linguistics), with over 100 peer-reviewed publications. He contributed to the organization of MathNLP at EMNLP 2022 and LREC-COLING 2024. André participated in 7 tutorials, and coorganized 1 conference and 6 workshops.

Causal Foundations for Interpretability: What Your Methods Actually Identify

Shruti Joshi, Aaron Mueller, Patrik Reizinger, Dhanya Sridhar

<https://causality4mechinterp.github.io/>

Interpretability findings on large language models often fail to generalise, and causal claims can outrun the evidence. A recurring source, we suggest, is a borrowed vocabulary: interpretability inherited the words of causality — "mechanisms", "circuits", "features" — without always inheriting the bookkeeping that keeps causal claims precise. This tutorial teaches that bookkeeping for practitioners. Participants leave able to take any interpretability result and answer three questions: what claim does this method license, what would strengthen the claim, and when will it generalise?

Shruti Joshi

email: shruti.joshi@mila.quebec

website: <https://shrutij01.github.io/>

Shruti Joshi is a PhD student at Mila and Université de Montréal, advised by Dhanya Sridhar. Her research asks how models can perform unsupervised discovery of physical mechanisms by collecting data to falsify their own hypotheses, and how to evaluate the representations they learn for the invariances they recover, drawing on identifiability theory and causal inference. She has extended this work to the interpretability and control of large language models, with a focus on safety through generalizable guarantees and on understanding models' reasoning processes.

Aaron Mueller

email: amueller@bu.edu

website: <https://aaronmueller.github.io/>

Aaron Mueller is an Assistant Professor of Computer Science, with a courtesy appointment in Data Science, at Boston University. His research addresses the robustness, efficiency, and interpretability of language models, spanning causal and mechanistic interpretability methods, linguistically informed evaluations of language models, and the development of more sample-efficient models. He completed his PhD at Johns Hopkins University's Center for Language and Speech Processing, advised by Tal Linzen and Mark Dredze, and held a Zuckerman postdoctoral fellowship at Northeastern University and the Technion.

Patrik Reizinger

email: patrik.reizinger@tuebingen.mpg.de

website: <https://rpatrik96.github.io>

Patrik Reizinger is an IVADO postdoctoral fellow in causal representation learning at Mila and Université de Montréal. His research works toward understanding deep learning through identifiable causal and self-supervised representation learning, out-of-distribution generalization, and compositionality in language models, extending machine learning theory to account for inductive biases grounded in physical and causal structure. He completed his doctoral research at the Max Planck Institute for Intelligent Systems in Tübingen and holds degrees in electrical engineering from the Budapest University of Technology and Economics, and is part of the ELLIS and IMPRS-IS programs.

Dhanya Sridhar

email: dhanya.sridhar@mila.quebec

website: <https://www.dsridhar.com/>

Dhanya Sridhar is an Assistant Professor in the Department of Informatics and Operations Research (DIRO) at Université de Montréal and a core academic member of Mila - Quebec Artificial Intelligence Institute, where she holds a Canada CIFAR AI Chair and co-leads the IVADO R3AI working group on safe and aligned AI. Her research combines causality and machine learning to build AI systems that are robust to distribution shifts, adapt efficiently to new tasks, and support scientific discovery — spanning causal representation learning under minimal supervision, causally grounded predictors, causal abstraction for interpreting large models, in-context learning, and responsible AI.

Multilingual and Multicultural LLMs

Wei Xu, David Ifeoluwa Adelani, Vinodkumar Prabhakaran, Sara Hooker, Alan Ritter

<https://emnlp2026-multilingual-cultural-llm.github.io/>

While large language models (LLMs) have substantially expanded support for non-English languages, their ability to generalize across typologically diverse languages and produce contextually appropriate, culturally grounded responses remains a significant challenge. This tutorial examines the rapidly evolving area of understanding, evaluating, and improving the multilingual capabilities and multicultural alignment of LLMs. We will primarily focus on cutting-edge advances and latest publications from 2025–2026, while incorporating selected introductory material to provide necessary background and to ensure accessibility for a broad audience. Participants will gain an up-to-date view of both methods and practical insights, including multilingual pre-, mid- and post-training, multilingual-specific optimization techniques, culturally aware evaluation, and key open problems in developing more fair, inclusive, and robust LLMs.

Wei Xu

email: wei.xu@cc.gatech.edu

website: <https://cocoxu.github.io/>

Wei Xu is an Associate Professor in the College of Computing and Machine Learning Center at the Georgia Institute of Technology. Her research interests include evaluation, robustness, multilingualism, cultural alignment, user simulation, and reasoning of LLMs. She is a recipient of the NSF CAREER Award, faculty research awards from Google, Sony, Crowdfunder, and Criteo, paper awards and honorable mentions at COLING’18, ACL’23, and ACL’24, including the Best Social Impact Paper Award at ACL 2024, titled “Having Beer after Prayer? Measuring Cultural Bias in Large Language Models.” She is confirmed to serve as a Program Chair for EMNLP 2027 and has regularly served as a senior area chair for *ACL conferences. She has given a tutorial on “Automatic and Human-AI Interactive Text Generation” at ACL 2024.

David Ifeoluwa Adelani

email: david.adelani@mila.quebec

website: <https://dadelani.github.io/>

David Ifeoluwa Adelani is an Assistant Professor at the McGill University School of Computer Science, a Core Academic Member at Mila - Quebec AI Institute, a Canada CIFAR AI Chair, and an IVADO Professor. His research interests include multilingual natural language processing and speech processing with a focus on low-resource languages. Some of his research on AfricaNLP have received Paper Awards including Best Paper Award (Global Challenges) at COLING 2022 for AfroX-LMR, a multilingual pre-trained language model for African languages, Area Chair Award at IJCNLP-AACL 2023 for MasakhaNEWS topic classification dataset, Outstanding Paper Award at NAACL 2025 for “IrokoBench—an LLM evaluation benchmark on African languages”, and the Best Theme Paper Award for “WorldCusines—a multicultural multilingual Visual QA dataset” at NAACL 2025. Previously, he gave a tutorial on “A Journey Through the Opportunity of Low Resourced Natural Language Processing — An African Lens” at NeurIPS 2021.

Vinodkumar Prabhakaran

email: vinodkpg@google.com

website: <https://cs.stanford.edu/~vinod/>

Vinodkumar Prabhakaran is a Senior Staff Research Scientist at Google Research, where he co-leads the interdisciplinary Technology, AI, Society and Culture (TASC) team. His research focuses on questions at the intersection of AI and society, with a focus on global and cross-cultural considerations. Before Google, he was a postdoc at Stanford University, and obtained his PhD from Columbia University. His prior research focused on building scalable ways using language technologies to identify and address large-scale societal issues such as racial disparities in policing, workplace incivility, and online abuse. He has published over 50 articles in top-tier venues such as the PNAS, ACL, TACL, NAACL, EMNLP, NeurIPS, and FAccT.

Sara Hooker

email: sara@adaptionlabs.ai

website: <https://www.sarahooker.me/>

Sara Hooker is a co-founder and CEO of Adaption Labs. She previously served as Vice President of Research at Cohere from 2022 to 2025, where she led Cohere Labs, a fundamental AI research lab that seeks to solve complex machine learning problems. Since 2026, she has been the Co-Founder and CEO of Adaptation Labs. Her research interests include model efficiency and AI accessibility, with a focus on multilingual, multimodal, model compression, ethical AI practices and fair LLM evaluation. She is a co-founder of the Trustworthy ML Initiative, serves on Kaggle’s ML Advisory Board and the World Economic Forum Council on the Future of Artificial Intelligence, and is a member of the MLC research group. Sara was also named one of TIME’s 100 Most Influential People in AI in 2024, and named one of AI’s “Top 13 Innovators” by Fortune in 2023. She received the Best Paper Award at ACL 2024, titled “Aya Model: An Instruction Finetuned Open-Access Multilingual Language Model”.

Alan Ritter

email: alan.ritter@cc.gatech.edu

website: <https://aritter.github.io/>

Alan Ritter is an Associate Professor in the College of Computing at the Georgia Institute of Technology. His research focuses on natural language processing and machine learning. He received his Ph.D. in Computer Science and Engineering from the University of Washington and subsequently completed a postdoctoral fellowship in the Machine Learning Department at Carnegie Mellon University. Alan served as Program Chair for NAACL 2025 and is the recipient of an NSF CAREER Award, an Amazon Research Award, a Sony Faculty Innovation Award, and a Google Academic Research Award, as well as Paper Awards at ACL 2023, 2024, and 2025. His research has been featured in media outlets including WIRED, TNW, Bloomberg, and VentureBeat.

Author Index

Adelani, David Ifeoluwa, 10

Cui, Hejie, 1

Cui, Yang, 2

Freitas, André, 6

Hofmann, Valentin, 4

Hooker, Sara, 10

Joshi, Shruti, 8

Kallini, Julie, 4

Kumar, Sachin, 4

Li, Manling, 1

Li, Zoey Sha, 1

Libovický, Jindřich, 4

Limisiewicz, Tomasz, 4

Meadows, Jordan, 6

Mueller, Aaron, 8

Prabhakaran, Vinodkumar, 10

Reizinger, Patrik, 8

Ritter, Alan, 10

Salesky, Elizabeth, 4

Sridhar, Dhanya, 8

Sun, Jingyuan, 2

Valentino, Marco, 6

Wang, Yifan, 2

Wang, Zihan, 1

Xu, Wei, 10

Zhang, Lan, 6

Zhang, Yunhao, 2