

EACL 2026

**Proceedings of the 19th Conference of the European Chapter
of the Association for Computational Linguistics**

Volume 6: Tutorial Abstracts

March 28-29, 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 979-8-89176-385-2

Introduction

Welcome to the Tutorials Session of EACL 2026!

The EACL tutorials session is designed to provide conference participants with comprehensive introductions to key topics in our rapidly evolving field, delivered by leading experts. These tutorials aim to foster learning, cross-fertilization of ideas, and engagement with emerging directions in computational linguistics and natural language processing.

This year, as in recent editions, the call, submissions, reviews, and tutorial selection were jointly coordinated for both EACL and ACL. A shared review committee was formed, comprising the EACL tutorial chairs (Chenghua Lin, Aline Paes, and Rodrigo Wilkens) and the ACL tutorial chairs (Jacob Andreas and Kenton Murray). Each proposal received at least two reviews. The selection process considered multiple criteria, including clarity, preparedness, novelty, timeliness, instructors' expertise, expected audience, commitment to open access teaching materials, diversity (in terms of multilingualism, gender, age, and geographic representation), and compatibility with venue preferences. In total, 60 tutorial proposals were submitted, of which 4 were selected for presentation at EACL 2026, covering introductory and cutting-edge research topics.

The tutorials featured in this volume address a diverse and timely set of themes at the intersection of language, artificial intelligence, and human cognition. They cover topics ranging from AI-assisted scientific discovery and experimentation to the cognitive and behavioral properties of large language models, including biases and reasoning patterns. The program also includes tutorials on computational approaches to modeling language in the brain, as well as on multimodal language models for human–AI interaction, encompassing agentic systems and inclusive applications.

We would like to thank all tutorial authors for their excellent contributions and for their flexibility throughout the process. We also express our sincere gratitude to the conference organizers for their effective collaboration, particularly to the General Chair, Aline Villavicencio, for her guidance and patience during the entire process.

We hope you enjoy the tutorials!

EACL 2026 Tutorial Co-chairs

Chenghua Lin

Aline Paes

Rodrigo Wilkens

Program

Saturday, March 28, 2026

- **09:00–12:30**
AI-assisted Scientific Discovery, Experimentation, Content Generation, and Evaluation
- **12:30–14:00**
Break
- **14:00–17:30**
Cognitive Effects and Biases in Large Language Models

Sunday, March 29, 2026

- **09:00–12:30**
Encoding and Decoding Language in the Brain with Language Models
- **09:00–12:30**
Multimodal Large Language Models for Human–AI Interaction: Foundations, Agents, and Inclusive Applications

Organizing Committee

General Chair

Aline Villavicencio, University of Exeter

Tutorial Chairs

Chenghua Lin, University of Manchester

Aline Paes, Universidade Federal Fluminense

Rodrigo Wilkens, University of Exeter

Table of Contents

<i>AI-assisted Scientific Discovery, Experimentation, Content Generation, and Evaluation</i>	
Yufang Hou, Steffen Eger, Anne Lauscher, Wei Zhao and Yong Cao	1
<i>Cognitive Effects and Biases in Large Language Models</i>	
Markus Schedl, Ralph Hertwig, Antonela Tommasel and Shahed Masoudian	4
<i>Encoding and Decoding Language in the Brain with Language Models</i>	
Anuja Negi, Mathis Lamarre, Christine Tseng and Subba Reddy Oota	7
<i>Multimodal Large Language Models for Human-AI Interaction: Foundations, Agents, and Inclusive Applications</i>	
Shafiq Joty, Enamul Hoque, Ahmed Masry, Spandana Gella and Samira Ebrahimi Kahou	9

T1: AI-assisted Scientific Discovery, Experimentation, Content Generation, and Evaluation

Yufang Hou, Steffen Eger, Anne Lauscher,
Wei Zhao, Yong Cao

<https://ai4science-tutorial.github.io/>

Cutting-edge

Pavillon DE RABAT (Level 1)

Saturday, March 28, 2026 - from 09:00 to 12:30

With the advent of large multimodal foundation models such as ChatGPT, Gemini, Claude, and DeepSeek, scientific research stands at the threshold of an AI-based technological transformation. Recent surveys indicate that the majority of researchers anticipate AI will become mainstream in scientific research within the next two years. This tutorial provides an in-depth overview of recent advances in AI-assisted tools and models that support and enhance the entire scientific research process, building upon findings from our recent survey paper. We will explore how AI is revolutionizing each stage of the research cycle: (1) Literature Search and Summarization – examining AI-enhanced search systems, paper chat interfaces, graph-based knowledge discovery, and personalized recommender systems; (2) Idea Generation and Experimentation – covering LLM-based hypothesis formulation, multi-agent systems, and automated experimentation tools; (3) Multimodal Content Understanding and Generation – surveying approaches to scientific figure comprehension, automatic diagram generation from text, and poster/slide creation; (4) Text-based Content and Table Generation – reviewing models for abstracts, citations, meta-analysis tables, and long-form content like survey papers; and (5) AI-supported Peer Review – introducing automated review analysis, feedback generation, and meta-review synthesis.

Yufang Hou, Interdisciplinary Transformation University, Austria

email: yufang.hou@it-u.at

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

Bio. Yufang Hou is a professor at IT:U – Interdisciplinary Transformation University Austria. At IT:U, she leads the NLP group with a strong focus on LLMs governance (with a particular focus on content veracity), computational argumentation, fact-checking, knowledge representation and reasoning, and human-centred NLP applications in education, science, and health. Yufang has served as an organizing committee member and senior program committee member for various NLP conferences, and has co-organized several workshops on topics such as argument mining, efficient NLP, and AI4Science.

Steffen Eger, University of Technology Nuremberg, Germany

email: steffen.eger@utn.de

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

Bio. Steffen Eger is a Full Professor at University of Technology Nuremberg, Germany, leading the Natural Language Learning & Generation. His research focuses on various aspects of Natural Language Processing (NLP), in particular evaluation for text generation and interdisciplinary problems at the intersection of NLP, the digital humanities and the social sciences. Recently, he has developed a keen interest in multimodal scientific content generation in the science domain. He has served the ACL community in various roles (SAC, AC, reviewer) and has extensive experience organizing ACL workshops in the context of evaluation and AI4Science.

Anne Lauscher, University of Hamburg, Germany

email: anne.lauscher@uni-hamburg.de

website: <https://anne-lauscher.de/>

Bio. Anne Lauscher is a Professor of Data Science at the University of Hamburg, where her research group investigates safe and responsible Natural Language Processing. Beyond her strong focus on ethical aspects of AI, she also developed an interest in mining scientific literature and citation analysis already during her Ph.D. studies. Further focusing on AI4Science, she recently presented works on analyzing peer reviews and supporting peer reviewers with targeted AI systems as well as on supporting physics experimentation through deploying LLMs in agentic setups. In the past, Anne has co-organized multiple workshops and served the ACL community in various roles, e.g., as senior area chair and as ethics reviewer.

Wei Zhao, University of Aberdeen, UK

email: wei.zhao@abdn.ac.uk

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

Bio. Wei Zhao is a Lecturer (=Assistant Professor) in NLP at the University of Aberdeen, Scotland. His research group works on evaluation of NLP systems, cross-temporal NLP, and more recently, AI for Science. He has served as an area chair and reviewer in the ACL community. He has also co-organized several NLP workshops in the ACL and COLM communities on the topics of evaluation of NLP systems and cross-temporal NLP, as well as AI4Science, the 1st Workshop on Human-LLM Collaboration for Ethical and Responsible Science Production, together with three other tutorial presenters.

Yong Cao, University of Tübingen, Germany

email: yong.cao@uni-tuebingen.de

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

Bio. Yong Cao is a postdoctoral researcher at the University of Tübingen, Germany. His research focuses on AI4Science, including multimodal document understanding, summarization, scientific impact prediction, and scientific document generation. Additionally, he also organizes workshops at ACL conferences and leads AI4Science Seminar events at the University of Tübingen. Yong's academic contributions have been featured in prestigious NLP conferences such as ACL, NAACL, EACL, ICLR, CoLM and NeurIPS including reviewer and area chair. He also serves on the program committee at AAAI 2022-2025.

T2: Cognitive Effects and Biases in Large Language Models

Markus Schedl, Ralph Hertwig

Antonela Tommasel, Shahed Masoudian

https://cpjku.github.io/cobis_llm_eacl2026/

Cutting-edge

Pavillon DE RABAT (Level 1)

Saturday, March 28, 2026 - from 14:00 to 17:30

Cognitive effects such as anchoring, positional effect, or confirmation bias are core aspects of human decision making and reasoning. As LLMs increasingly act as communicative partners, reasoning tools, and evaluators, understanding how these cognitive effects influence their behavior and vice versa has become essential. While recent studies have adapted psychological experiments to detect cognitive biases in LLMs, they often use a particular kind of experimental setup from psychology that carries implications for human performance. In addition, current NLP studies often confuse cognitive effects with biases, diverging from their psychological foundations and overlooking potentially functional aspects of these phenomena.

In this tutorial, jointly organized by NLP researchers and a cognitive psychologist and decision scientist, we aim to build shared conceptual and methodological ground between the two disciplines. We begin by outlining how cognitive effects and biases are defined, validated, and sometimes debated within psychology, highlighting differences and contradictions in experimental designs. We then bridge these insights to NLP through an overview of key studies examining cognitive biases in LLMs, mapping their methodological parallels and divergences. The tutorial also includes a hands-on component where participants explore the challenges of detecting a single cognitive bias (e.g., positional bias) in multilingual LLMs, illustrating the nuances and pitfalls of such evaluations. We conclude by discussing emerging research directions and open questions at the intersection of cognitive science and large language models.

Markus Schedl, Johannes Kepler University Linz (JKU), Germany

email: markus.schedl@jku.at

website: <http://www.mschedl.eu>

Bio. Markus Schedl is a Full Professor at the Johannes Kepler University Linz (JKU), affiliated with the Institute of Computational Perception, leading the Multimedia Mining and Search Group and the Human-centered AI Group at the Linz Institute of Technology (LIT) AI Lab. His research interests include information retrieval, recommender systems, natural language processing, and trustworthy AI, with a keen interest in psychological aspects of AI-powered systems and their users. He (co-)authored more than 300 refereed conference papers, journal articles, books, and book chapters, and has given numerous tutorials at top information retrieval and recommender systems venues.

Ralph Hertwig, Center of Adaptive Rationality, Max Planck Institute for Human Development, Germany

email: hertwig@mpib-berlin.mpg.de

website: <https://www.mpib-berlin.mpg.de/staff/ralph-hertwig>

Bio. Ralph Hertwig is the Director of the Center of Adaptive Rationality (ARC) at the Max Planck Institute for Human Development in Berlin. He is also an Honorary Professor at the Free University Berlin and the Humboldt University Berlin. His research focuses on how people make decisions and how good decisions can be fostered. He is a member of the German Academy of Sciences Leopoldina, the Berlin-Brandenburg Academy of Sciences and Humanities, and the acatech, the National Academy of Science and Engineering. He has received the Leibniz Prize, Germany's most prestigious research award.

Antonela Tommasel, Johannes Kepler University Linz, Germany / UNCPBA, Argentina

email: antonela.tommasel@jku.at

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

Bio. Antonela Tommasel is an Assistant Professor at the Johannes Kepler University Linz (JKU) and at the National University of the Center of the Province of Buenos Aires (UNCPBA, Argentina), as well as a researcher at the National Scientific and Technical Research Council (CONICET, Argentina). Her research lies at the intersection of artificial intelligence and social computing, with a focus on recommender systems. Her work explores fairness and bias in LLM-based recommendations, agentic LLM systems for supporting complex decision-making in software engineering, and the dynamics of user engagement and misinformation in online platforms.

Shahed Masoudian, Johannes Kepler University Linz (JKU), Germany

email: shahed.masoudian@jku.at

website: <https://hcai.at/persons/masoudian>

Bio. Shahed Masoudian is a University Assistant and PhD student at the Johannes Kepler University Linz (JKU). A core part of his work focuses on the controllability of language models, treating user control over model behavior as a practical and promising solution to addressing biases within language models. He develops novel methods that allow bias mitigation to be dynamic, scalable, and adaptable to different contexts, without requiring extensive retraining. He also supervises students on various topics in information retrieval, recommender systems, controllable text generation, and fairness of language models.

T3: Encoding and Decoding Language in the Brain with Language Models

Anuja Negim, Mathis Lamarre,
Christine Tseng, Subba Reddy Oota

<https://subbareddy248.github.io/tutorials/eacl-2026-brain-ai/>

Introductory

Pavillon DE RABAT (Level 1)

Sunday, March 29, 2026 - from 09:00 to 12:30

This EACL 2026 tutorial will cover the foundations of brain–language model alignment and will then explore recent advances on scaling laws of language models for brain alignment, multilingual brain encoding, recent developments in fine-tuning language models with brain data, and advances in brain decoding using language models, including semantic reconstruction of continuous language from brain data. Participants will gain an overview of current naturalistic datasets, computational frameworks, and methods driving the emerging field of NeuroAI. The learning objectives are: (1) Understand the fundamental concepts of brain-AI alignment and encoding models; (2) Learn methodologies for comparing brain activity with model representations; (3) Gain insights into multilingual processing in both human brains and language models; (4) Master techniques for brain-informed model fine-tuning and evaluation; and (5) Discover practical applications and future research directions in Neuro-AI.

Anuja Negi, TU Berlin, Germany

email: anuja.negi@tu-berlin.de

website: <http://anujanegi.me/>

Bio. Anuja Negi is a PhD candidate at the Cognitive Computing in Biological and Artificial Systems Lab at TU Berlin, Germany is affiliated with the Bernstein Center for Computational Neuroscience (BCCN) Berlin. She received her Master’s degree in Computational Neuroscience jointly from Technische Universität Berlin and Humboldt-Universität zu Berlin. Her research focuses on language comprehension and learning, bridging neuroscience and artificial intelligence, using insights from one to better understand and improve the other. Her work has appeared at Conference on Neural Information Processing Systems (NeurIPS), Society for

Neuroscience (SfN), Cognitive Computational Neuroscience (CCN), the Society for the Neurobiology of Language (SNL), Organization for Human Brain Mapping (OHBM), and NeuroImage journal.

Mathis Lamarre, TU Berlin, Germany

email: m.lamarre@campus.tu-berlin.de

website: <https://github.com/Matlmr>

Bio. Mathis Lamarre is a PhD candidate at the Cognitive Computing in Biological and Artificial Systems Lab at TU Berlin, Germany is affiliated with the Bernstein Center for Computational Neuroscience (BCCN) Berlin. He received his Master's degree in Computational Science and Engineering from ETH Zürich, and his Bachelor's degree in Life Sciences and Technology from EPF Lausanne. He is interested in using tools from Natural Language Processing to better understand language processing in the brain, in particular across languages. His work has appeared at the Organization for Human Brain Mapping (OHBM), Cognitive Computational Neuroscience (CCN) and Empirical Methods in Natural Language Processing (EMNLP).

Christine Tseng, TU Berlin, Germany

email: tseng@tu-berlin.de

website: <https://www.linkedin.com/in/ctseng12>

Bio. Dr. Christine Tseng is a Postdoctoral Researcher at TU Berlin, Germany, and completed their Ph.D. at Helen Wills Neuroscience Institute at UC Berkeley, USA. Their research interests include language processing in the brain, social neuroscience, brain encoding models, and using AI/ML to better understand the brain. Their work has appeared at CogSci, Society for Neuroscience (SfN), Cognitive Computational Neuroscience (CCN), the Society for the Neurobiology of Language (SNL), and in the Journal of Neuroscience.

Subba Reddy Oota, TU Berlin, Germany

email: subba.reddy.oota@tu-berlin.de

website: <https://subbareddy248.github.io/>

Bio. Dr. Subba Reddy Oota is a Postdoctoral Researcher at TU Berlin, Germany, and completed his Ph.D. at Inria Bordeaux, France. His research focuses on language analysis in the brain, brain encoding/decoding, multimodal information processing, and interpreting AI models. He has published at top venues including NeurIPS, ICLR, ACL, EMNLP, NAACL, INTERSPEECH, and TMLR.

T4: Multimodal Large Language Models for Human–AI Interaction: Foundations, Agents, and Inclusive Applications

Shafiq Joty, Enamul Hoque, Ahmed Masry,
Spandana Gella, Samira Ebrahimi Kahou

<https://mllm4haii.github.io/>

Cutting-edge

Salle Le Riad

Sunday, March 29, 2026 - from 09:00 to 12:30

Multimodal large language models (MLLMs) are redefining how humans communicate and collaborate with machines. They extend the capabilities of text-based LLMs to perceive, reason, and act across text, images, charts, forms, and graphical user interfaces (GUIs). These models are now capable of answering questions about charts, summarizing infographics, operating software through natural language, and supporting multilingual and accessible visualization.

This tutorial offers a concise, three-hour introduction to the foundations, agentic capabilities, and inclusive applications of MLLMs, with a focus on visually grounded and interactive language tasks. We will cover core architectural designs (encoders, connectors, fusion and decoding mechanisms), multimodal alignment and learning strategies, and reasoning techniques for structured visuals such as charts, forms, and infographics. The tutorial then examines multimodal and conversational agents that perform dialogue-driven reasoning and co-creative analysis in graphical user interfaces. We conclude with discussions on accessibility, multi-lingual communication, responsible deployment, and future challenges in building human-centered multimodal AI.

Shafiq Joty, Salesforce Research, USA

email: sjoty@salesforce.com

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

Bio. Shafiq Joty is a Research Director at Salesforce Research, and is also an Associate Professor (on leave) at NTU, Singapore. His work has primarily focused on developing language analysis tools and NLP applications. A significant part of his current research focuses on multilingual (machine translation, cross-lingual

transfer), multimodal (visual-language learning, NLP+Vis, Code+NLP) NLP, interpretability and robustness of NLP models. His research contributed to 17 patents and more than 110 papers in top-tier NLP and ML conferences and journals including ACL, EMNLP, NAACL, NeurIPS, ICML, ICLR, CVPR, ECCV, ICCV, CL and JAIR. Shafiq served (or will serve) as a PC chair of SIGDIAL'23, an S/AC for ICLR-23, ACL'22, EMNLP'21, ACL'19-21, EMNLP'19, NAACL'21 and EACL'21 and an AE for ACL-RR. He gave tutorials at IEEE Vis'22, ACL'19, ICDM'18 and COLING'18, and taught deep learning for NLP,¹ a graduate-level NLP course, and an undergraduate NLP course at NTU.

Enamul Hoque, York University, Canada

email: enamulh@yorku.ca

website: <https://www.yorku.ca/enamulh/>

Bio. Enamul Hoque is an Associate Professor at York University where he directs the Intelligent Visualization Lab. Previously, he was a postdoctoral fellow in Computer Science at Stanford University. He received the Ph.D. degree in Computer Science from the University of British Columbia. His research focuses on combining information visualization and human-computer interaction with natural language processing to address the challenges of the information overload problem. Recently, he has worked on developing natural language interfaces for visualizations, automatic chart question answering, chart retrieval and chart summarization. He has also worked on developing visual text analytics to support the user's task of exploring and analyzing conversations. Since his research is uniquely positioned at the intersection of information visualization, NLP, and HCI, he publishes at the major venues in each of these areas such as IEEE Vis, ACL, EMNLP, CHI, and UIST. He serves as an Area Chair for the ACL Rolling Review (2021-) and as a program committee member (2018-) for the IEEE Vis.

Ahmed Masry, York University, Canada

email: masry20@yorku.ca

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

Bio. Ahmed Masry is a PhD student at York University, Canada, supervised by Professor Enamul Hoque. He previously interned at ServiceNow Research and Mila. His research focuses on developing benchmarks and vision-language models for chart and document understanding, with an emphasis on supervised fine-tuning datasets, reinforcement learning, and improving vision-language alignment in multimodal architectures. Ahmed has led and contributed to popular benchmarks like ChartQA, ChartQPro and Chart-to-Text, as well as models includ-

¹https://ntunlp.sg.github.io/ce7455_deep-nlp-20/

ing ChartGemma, BigCharts-R1, and AlignVLM. His work has been published in leading NLP and ML venues like ACL, EMNLP, COLM, NeurIPS, ICLR. He also serves as a reviewer for ACL Rolling Review, NeurIPS, and ICLR. Ahmed has received several prestigious Canadian national and provincial awards, including the NSERC CGS-D, Ontario Graduate Scholarship, and Mitacs Accelerate Award. His research was also recognized with the Best Paper Award at the ChartQA Workshop at CVPR 2021.

Spandana Gella, ServiceNow Research, Canada

email: spandana.gella@servicenow.com

website: <https://www.servicenow.com/research/author/spandana-gella.html>

Bio. Spandana Gella is a Research Scientist at ServiceNow Research. She holds a Ph.D. in Computer Science from the University of Edinburgh. Her research focuses on building robust and safe frontier models and autonomous LLM-agents. In the past, she co-organized multiple workshops co-located with top-tier conferences including Representation Learning for NLP (2018, 2019, 2020), Shortcomings in Vision and Language (2018, 2019), and the Workshop on Multilingual Multimodal Learning (2022).

Samira Ebrahimi Kahou, University of Calgary, Canada

email: samira.ebrahimikahou@ucalgary.ca

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

Bio. Samira Ebrahimi Kahou is an Associate Professor at the Electrical and Software Engineering Department at the University of Calgary and an adjunct professor at the School of Computer Science at McGill University. She holds a CIFAR AI Chair. Ebrahimi Kahou’s pioneering work in visual reasoning includes the two well-known datasets “Something Something” and “FigureQA”. Ebrahimi Kahou and her group currently work on solving fundamental problems in representation learning for decision making, with a broad focus on generalization and efficient learning. Besides this primary focus, she also has expertise in knowledge distillation, climate modeling using deep learning, building large-scale datasets, clinical decision making, and NLP. Her publications appear in leading venues such as NeurIPS, ICLR, ICML, CVPR, and ICCV, and she has served as an Area Chair for NeurIPS, ICCV, and EMNLP.

Author Index

Cao, Yong, 1

Eger, Steffen, 1

Gella, Spandana, 9

Hertwig, Ralph, 4

Hoque, Enamul, 9

Hou, Yufang, 1

Joty, Shafiq, 9

Kahou, Samira Ebrahimi, 9

Lamarre, Mathis, 7

Lauscher, Anne, 1

Masoudian, Shahed, 4

Masry, Ahmed, 9

Negi, Anuja, 7

Oota, Subba Reddy, 7

Schedl, Markus, 4

Tommasel, Antonela, 4

Tseng, Christine, 7

Zhao, Wei, 1