

ACL 2025

**The 63rd Annual Meeting of the Association for  
Computational Linguistics**

**Tutorial Abstracts**

July 27, 2025

©2025 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](mailto:acl@aclweb.org)

ISBN 979-8-89176-255-8

## Introduction

Welcome to the tutorial session of ACL 2025!

We are thrilled to welcome you to this year’s ACL tutorial session, a cornerstone event of the conference. Our tutorials offer a deep dive into a diverse range of cutting-edge and emerging topics, presented by leading researchers at the forefront of their fields. These sessions are designed to equip you with the latest insights, tools, and methodologies, enriching your understanding of the ever-evolving landscape of computational linguistics and natural language processing.

Continuing our longstanding tradition, the call, submission, review, and selection processes for tutorials were jointly coordinated across NAACL, ACL, and EMNLP. The review committee comprised the NAACL tutorial chairs (Rui Zhang, Swabha Swayamdipta, and Maria Lomeli), the ACL tutorial chairs (Yuki Arase, David Jurgens, and Fei Xia), and EMNLP tutorial representatives (Jessy Li and Fei Liu). Each tutorial proposal underwent a thorough evaluation by a panel of two to three reviewers, who assessed the submissions based on several key criteria: clarity and preparedness, novelty and timeliness of the topic, instructors’ experience, anticipated audience interest, commitment to open access teaching materials, and diversity considerations—including multilingualism, gender, age, and geographic representation—as well as compatibility with the preferred conference venue. The tutorial chairs for all three conferences met to discuss each tutorial and make the final decisions and assignments to each conference.

This year, we received a remarkable total of 76 tutorial submissions, many of which were very engaging and made for a highly competitive selection process. Ultimately, 8 tutorials were selected for presentation at ACL. These tutorials promise to offer engaging and informative sessions, catering to a broad spectrum of interests and expertise levels within the computational linguistics and NLP community.

We deeply appreciate the contributions of all tutorial authors, as well as the tireless efforts and strong collaboration of the conference organizers—with special thanks to General Chair Roberto Navigli for his exceptional leadership and Jenn Rachford for her work in scheduling and coordinating tutorials at the conference.

We hope these tutorials provide an exciting and engaging start to your overall experience at ACL 2025. Wishing you an inspiring and engaging time at the conference, filled with learning and connection!

Warm regards,  
ACL 2025 Tutorial Co-Chairs,  
Yuki Arase, David Jurgens, and Fei Xia

# Organizing Committee

## General Chair

Roberto Navigli, Sapienza University Rome & Babelscape

## Program Chairs

Wanxiang Che, Harbin Institute of Technology

Joyce Nabende, Makerere University

Ekaterina Shutova, University of Amsterdam

Mohammad Taher Pilehvar, Cardiff University / Tehran Institute for Advanced Studies

## Tutorial Chairs

Yuki Arase, Tokyo Institute of Technology

David Jurgens, University of Michigan

Fei Xia, University of Washington

## Table of Contents

<i>Inverse Reinforcement Learning Meets Large Language Model Alignment</i> Mihaela van der Schaar and Hao Sun .....	1
<i>Eye Tracking and NLP</i> David Reich, Omer Shubi, Lena Jäger and Yevgeni Berzak .....	2
<i>Uncertainty Quantification for Large Language Models</i> Artem Shelmanov, Maxim Panov, Roman Vashurin, Artem Vazhentsev, Ekaterina Fadeeva and Timothy Baldwin .....	3
<i>Human-AI Collaboration: How AIs Augment Human Teammates</i> Sherry Wu, Diyi Yang, Joseph Chang, Marti A. Hearst and Kyle Lo .....	5
<i>Navigating Ethical Challenges in NLP: Hands-on strategies for students and researchers</i> Luciana Benotti, Fanny Ducel, Karën Fort, Guido Ivetta, Zhijing Jin, Min-Yen Kan, Seunghun J. Lee, Minzhi Li, Margot Mieskes and Adriana Pagano .....	7
<i>NLP for Counterspeech against Hate and Misinformation (CSHAM)</i> Daniel Russo, Helena Bonaldi, Yi-Ling Chung, Gavin Abercrombie and Marco Guerini .....	9
<i>Synthetic Data in the Era of Large Language Models</i> Vijay Viswanathan, Xiang Yue, Alisa Liu, Yizhong Wang and Graham Neubig .....	11
<i>Guardrails and Security for LLMs: Safe, Secure and Controllable Steering of LLM Applications</i> Traian Rebedea, Leon Derczynski, Shaona Ghosh, Makesh Narsimhan Sreedhar, Faeze Brahmaan, Liwei Jiang, Bo Li, Yulia Tsvetkov, Christopher Parisien and Yejin Choi .....	13

# Program

## Sunday, July 27, 2025

- 09:00 - 12:30     *Inverse Reinforcement Learning Meets Large Language Model Alignment*
- 09:00 - 12:30     *Eye Tracking and NLP*
- 09:00 - 12:30     *Uncertainty Quantification for Large Language Models*
- 09:00 - 12:30     *Human-AI Collaboration: How AIs Augment Human Teammates*
- 14:00 - 12:30     *Break*
- 14:00 - 17:30     *Navigating Ethical Challenges in NLP: Hands-on strategies for students and researchers*
- 14:00 - 17:30     *NLP for Counterspeech against Hate and Misinformation*
- 14:00 - 17:30     *Synthetic Data in the Era of Large Language Models*
- 14:00 - 17:30     *Guardrails and Security for LLMs: Safe, Secure, and Controllable Steering of LLM Applications*