

Contents

Table of Contents	i
1 Conference Information	1
General Chair Message	1
Program Chair Message	4
ACL Business and Event Manager Message	8
Organizing Committee	9
Program Committee	12
2 Anti-harassment Policy	17
Ethics Policy	17
3 Meal Info	19
4 Social Events	21
ACL 2025 Welcome Reception	21
ACL 2025 Social Event	22
5 Keynotes	23
Monday, July 28	23
Tuesday, July 29	25
6 Panel Discussion	27
Monday, July 28	27
7 Tutorials	31
Overview	31
Tutorial 1	33
Tutorial 2	34
Tutorial 3	35

Tutorial 4	36
Tutorial 5	37
Tutorial 6	38
Tutorial 7	39
Tutorial 8	40
8 Main Conference	41
Main Conference: Monday, July 28, 2025	41
Main Conference: Tuesday, July 29, 2025	43
Main Conference: Wednesday, July 30, 2025	45
9 Workshops	47
Overview	47
W01: 6th AfricaNLP Workshop: Multilingual and Multicultural-aware LLMs	49
W02: The 12th Workshop on Argument Mining	50
W03: 2nd Workshop on Natural Language Processing meets Climate Change	51
W04: Eighth Workshop on Fact Extraction and VERification (FEVER)	52
W05: Workshop on Meaningful, Efficient, and Robust Evaluation of LLMs and GEM: Natural Language Generation, Evaluation, and Metrics	53
W06: The 19th Linguistic Annotation Workshop (LAW XIX)	54
W07: The 4th Workshop on NLP for Positive Impact	55
W08: REALM: First Workshop for Research on Agent Language Models	56
W09: SDP 2025: The 5th Workshop on Scholarly Document Processing	57
W10: The Third Workshop on Social Influence in Conversations (SICon 2025)	58
W11: Slavic NLP-2025: The 10th Biennial Workshop on Slavic NLP	59
W12: The 4th Table Representation Learning Workshop at ACL 2025	60
W13: The Fourth Ukrainian Natural Language Processing Workshop (UNLP 2025)	61
W14: The 20th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2025)	62
W15: 29th Conference on Computational Natural Language Learning (CoNLL)	63
W16: The 22nd International Conference on Spoken Language Translation (IWSLT 2025)	64
W17: The 19th Workshop on Semantic Evaluation (Semeval-2025)	65
W18: The Second Workshop on Analogical Abstraction in Cognition, Perception, and Language (Analogy-ANGLE II)	66
W19: BioNLP 2025 and Shared Tasks (BioNLP-ST 2025)	67
W20: The 6th Workshop on Gender Bias in Natural Language Processing (GeBNLP 2025)	68
W21: The First Workshop on Large Language Model Memorization (L2M2)	69
W22: First Workshop on LLM Security (LLMSEC)	70
W23: MAGMaR: Multimodal Augmented Generation via Multimodal Retrieval	71
W24: FieldMatters–SIGTYP: The 7th Workshop on Research in Computational Linguistic Typology and Multilingual NLP (SIGTYP) and the Fourth Workshop on NLP Applications to Field Linguistics (FieldMatters)	72
W25: The 3rd workshop on Towards Knowledgeable Foundation Models	73
W26: WikiNLP: Advancing Natural Language Processing for Wikipedia	74
W27: The 9th Workshop on Online Abuse and Harms (WOAH 2025)	75
W28: The First Joint Workshop on Large Language Models and Structure Modeling	76
10 Poster Printing	77
11 Local Guide	79
12 Venue Map	83

Conference Information

General Chair Message

It is my great pleasure and honor to welcome you to the 63rd Annual Meeting of the Association for Computational Linguistics (ACL 2025), held in beautiful Vienna, Austria, from July 27 to August 1, 2025. ACL 2025 continues our field’s tradition of excellence in scholarship, innovation, and inclusivity, and I am deeply grateful to the many volunteers who have worked tirelessly to bring this event to life.

I want to express my deepest thanks to our Program Chairs — **Wanxiang Che**, **Joyce Nakatumba-Nabende**, **Mohammad Taher Pilehvar** and **Ekaterina Shutova** — who have overseen the reviewing and selection process, and shaped a compelling and diverse scientific program. Overall, we received more than 8300 submissions and accepted 1700 at the main conference and 1392 as findings. This was made possible through close coordination with the ACL Rolling Review (ARR) team, led by the ARR Editors-in-Chief — **Jun Suzuki**, **Jing Jiang**, and **Xiaodan Zhu** — and I sincerely thank them. Although some improvements are still possible, the ARR workflow was already well integrated with the conference review and decision-making process. We are also deeply grateful to the (many!) Senior Area Chairs, Area Chairs, reviewers, and the Best Paper Committee (led by **Rada Mihalcea** and **Roi Reichart**), whose dedication ensured the high quality of our program. I am also grateful to our Technical Open Review Chairs, **Niket Tandon** and **Lizhen Qu**, for their behind-the-scenes work. For those joining virtually, our Virtual Infrastructure Chairs — **Manling Li**, **Yang Liu**, and **Avi Sil** — have worked to make the hybrid conference experience inclusive and engaging. I would also like to recognize the efforts of our Ethics Chairs — **Karèn Fort** and **Bjorn Ross** — who have led the work aimed at ensuring that submissions meet the ethical standards of our field.

ACL 2025 hosts a rich set of workshops and tutorials, thanks to the dedicated efforts of our Workshop Chairs — **Terra Blevins** and **Christophe Gravier** — and our Tutorial Chairs — **Yuki Arase**, **David Jurgens**, and **Fei Xia**. Our Demonstration Track Chairs — **Pushkar Mishra**, **Smaranda Muresan**, and **Tao Yu** — have put together an impressive set of system demonstrations of cutting-edge innovations in NLP. I am proud to say that we also have a rich Industry Track, thanks to **Yunyao Li** and **Georg Rehm** who created a “conference in the conference”. Given the growing interaction between academia and industry, I hope this tradition, which stems from NAACL and which I reinstated at ACL, will be continued in the following editions.

As usual, TACL and CL papers are presented at the conference. Thanks go to the TACL Editors-in-Chief

(**Asli Celikyilmaz, Roi Reichart, and Dilek Hakkani-Tur**) and the CL Editor-in-Chief, **Wei Lu**, for their coordination efforts. I am also proud to announce that this year we are introducing the new CL Doctoral Dissertation Award at the conference. Other notable innovations include virtual presentation sessions held in parallel with the in-person poster sessions, a reception event dedicated to Findings presentations, and the introduction of combined oral and panel sessions designed to foster more thorough and interactive research discussions.

This conference could only have been made possible thanks to the invaluable work of our ACL Business Manager and Director of Events, **Jenn Rachford**, and her team — in particular, **Megan Haddad**, our ACL Administrator and Registration Manager — especially given the largest number of participants ever at ACL, estimated at the time of writing to be around 5000! Thanks also go to **Damira Mršić** and the Underline team for managing all content for the hybrid experience. The Local Organization Chairs, **Benjamin Roth** and **Dagmar Gromann**, worked hard on the bidding process, providing huge support with all local arrangements, the social event, and the visa process, eventually ensuring an outstanding experience for our in-person attendees. Our Visa Chairs — **Rexhina Blloshmi** and **Eleni Ilkou** — have helped numerous participants with their visa needs and any logistical difficulties in attending the conference. I also thank our Student Volunteer Chairs — **Pedro Henrique Luz de Araujo** and **Eleonora Mancini** — for their behind-the-scenes work, which will prove indispensable during the conference.

Additional thanks go to **Zhu Liu, Mingyang Wang, and Jin Zhao**, our Student Research Workshop Chairs, who, with the support of faculty advisors **Lea Frermann, Daniel Hershcovich, and Tristan Miller**, devoted their work to supporting and mentoring the next generation of researchers. They also secured additional funding from the Vienna Meeting Fund — congratulations!

Our Publication Chairs — **Pierpaolo Basile, Libo Qin, and Zhenghao Liu** — have ensured timely and high-quality conference proceedings. I also thank the Handbook Chairs, **Els Lefever and Qionikai Xu**, for assembling the handbook that will guide attendees through the conference program.

Of course our conference could not exist without the support of our sponsors! We are deeply grateful to our Sponsorship Chairs — **Raffaella Bernardi and Thomas Scialom** — and to **Chris Callison-Burch**, ACL's Sponsorship Director, for securing the generous support that enables us to keep the conference accessible to as many participants as possible.

Fostering an inclusive environment remains a key goal of ACL. Our Diversity and Inclusion Chairs — **Senja Pollak, Maria Ryskina, Shane Storks, and Hwaran Lee** — have worked to support diverse participation and organize activities that reflect the global nature of our vibrant community. An inclusive ACL allows researchers from all areas and backgrounds to contribute, collaborate, and thrive — enriching both the scientific conversation and the broader societal impact of our work. We remain committed to building a community where everyone feels welcomed and valued.

Thank you to the Publicity and Social Media Chairs — **Anette Frank, Shruti Rijhwani, and Horacio Saggion** — for their communication work through social media, extending the conference's reach to a wider audience. And this brings me to our Website and App Chairs — **Xudong Han and Alessandro Raganato** — who kept our community informed with timely updates about all aspects of the conference. The huge number of emails and communications that we received was smoothly managed throughout the year by our Internal Communication Chairs, **Sara Tonelli** and **Yiquan Wu**. We are very grateful!

I am particularly proud of the creation of a new role, the Documentation Chair, taken by **Chenghua Lin**, who did strenuous work to assemble a new handbook with up-to-date information coming from all chairs — whom I also thank for their valuable contributions — and maintain a clear and consistent record of all our processes. Given the speed at which all conference roles evolve in their duties and commitments, I hope this role will become standard in upcoming conferences to ensure a smooth transition.

Last but not least, I wish to thank the ACL Executive Committee for their guidance, and my fellow past General and Program Chairs, and especially **Claire Gardent** and **Colin Cherry**, for generously sharing their experience and insights.

ACL 2025 is a community-wide effort. Whether you are presenting a paper, leading a workshop or tutorial, volunteering, mentoring, supporting or attending to learn and connect, your presence and participation make a difference at this conference. Thank you for being part of it!

Welcome to Vienna. Welcome to ACL 2025!

Roberto Navigli

Sapienza University of Rome & Babelscape
General Chair, ACL 2025

Program Chair Message

Welcome to the 63rd Annual Meeting of the Association for Computational Linguistics!

ACL 2025 will be held in a hybrid format, offering attendees the option to join us in person in vibrant Vienna, Austria, or to participate remotely from anywhere in the world. Organizing ACL 2025 has been a collaborative effort, made possible by the dedication and hard work of thousands of people. We gratefully acknowledge the support and contributions of the following people:

- The General Chair, Roberto Navigli;
- The ARR Editors-in-Chief of the February 2025 cycle (Jun Suzuki, Jing Jiang, and Xiaodan Zhu) and the entire team (Mausam, Viviane Moreira, Vincent Ng, Lilja Øvrelid, Anna Rogers, Michael White, Margot Mieskes, Sarvnaz Karimi);
- The technical OpenReview chairs, Niket Tandon, Lizhen Qu, and the OpenReview support team, in particular Rachel, for multiple rounds of technical help in setting up ACL 2025 on the Open Review platform;
- The 169 Senior Area Chairs;
- The 1,937 Area Chairs and the 11,720 reviewers;
- The best paper committee chairs, Rada Mihalcea and Roi Reichart, and the best paper committee members;
- The ethics chairs, Karën Fort and Bjorn Ross;
- The workshop chairs, Terra Blevins and Christophe Gravier;
- The tutorial chairs, Yuki Arase, David Jurgens and Fei Xia;
- The industry track chairs, Yunyao Li and Georg Rehm;
- The demonstration chairs, Pushkar Mishra, Smaranda Muresan, and Tao Yu;
- The internal communications chairs, Sara Tonelli and Yiquan Wu;
- The website and conference app chairs, Xudong Han and Alessandro Raganato;
- The publication chairs, Pierpaolo Basile, Libo Qin, and Zhenghao Liu;
- The handbook chairs, Els Lefever and Qionghai Xu;
- The local organization chairs, Benjamin Roth and Dagmar Gromann, and their team;
- The visa chairs, Rexhina Blloshmi and Eleni Ilkou;
- The publicity and social media chairs, Anette Frank, Shruti Rijhwani and Horacio Saggion;
- The documentation chair, Chenghua Lin;
- The student research workshop chairs, Zhu Liu, Mingyang Wang and Jin Zhao;
- The student research workshop chairs faculty advisors, Lea Frermann, Daniel Hershcovich and Tristan Miller;

-
- The student volunteer chairs, Pedro Henrique Luz de Araujo and Eleonora Mancini;
 - The diversity and inclusion chairs, Senja Pollak, Maria Ryskina, Shane Storks and Hwaran Lee;
 - The sponsorship chairs, Raffaella Bernardi and Thomas Scialom;
 - The virtual infrastructure chairs, Manling Li, Yang Liu and Avi Sil;
 - The ACL Anthology Director Matt Post and his team;
 - The TACL editors-in-chief (Asli Celikyilmaz, Roi Reichart, Dilek Hakkani-Tur) and CL Editor in-Chief Wei Lu for coordinating TACL and CL presentations with us;
 - The ACL 2024 Program Chairs, Lun-Wei Ku, André F. T. Martins, Vivek Srikumar, for information and support;
 - Damira Mrsic and Underline Team;
 - Jennifer Rachford and entire conference support staff;
 - All the authors of papers who submitted their papers for review in the ARR 2025 February cycle and those who committed to the ACL 2025 conference.

Review Process All submissions to ACL 2025 went through a two-stage review process. First, papers were submitted to the ACL Rolling Review (ARR), where they were reviewed by reviewers and received meta-reviews from Area Chairs. Then, authors had the option to commit their reviewed papers to ACL via a separate ACL 2025 commitment site. At this stage, Senior Area Chairs provided recommendations, and final acceptance decisions were made by the Program Chairs. This process is consistent with previous conferences, including ACL 2024, EACL 2024, and NAACL 2025.

We worked closely with the ARR team, particularly the ARR February 2025 Editors-in-Chief, and served as guest Editors-in-Chief for this round. We helped recruit new reviewers and Area Chairs to ARR, resulting in 11,720 reviewers and 1,942 Area Chairs in the February 2025 ARR cycle to which most ACL 2025 papers were submitted. ACL also recruited 169 Senior Area Chairs to oversee the review and meta-review process. Overall, the ARR process ran smoothly, ensuring that all submitted papers received at least three reviews and a meta-review. For the ACL commitment phase, Senior Area Chairs made recommendations for 5,356 committed papers based on the reviews, meta-reviews, and the papers themselves, with final acceptance decisions made by the Program Chairs.

Acceptance Rate

In total, there are 1,699 papers accepted to the Main Conference and 1,392 papers accepted to Findings. The acceptance rate calculation follows precedent set by previous conferences that go through ACL Rolling Review (ARR), e.g. ACL 2024. The calculation takes into account the multi-stage process of ARR where a paper may get revised in ARR and then later committed to the conference. In total, we had 8,360 unique submissions across the December 2024 and February 2025 ARR cycles of which 5,501 papers were committed to ACL. The acceptance rate is 20.3% for the Main Conference papers and a further 16.7% for Findings papers.

Special Theme: Generalization of NLP Models

ACL 2025’s special theme is generalization of NLP models. Generalization is crucial for ensuring that models behave robustly, reliably, and fairly when making predictions on data different from their training data. Achieving good generalization is critically important for models used in real-world applications, as they should emulate human-like behavior. Humans are known for their ability to generalize well, and models should aspire to this standard. The theme track invites empirical and theoretical research and position and survey papers reflecting on the Generalization of NLP Models. The possible topics of discussion

include (but are not limited to) the following: 1. How can we enhance the generalization of NLP models across various dimensions—compositional, structural, cross-task, cross-lingual, cross-domain, and robustness? 2. What factors affect the generalization of NLP models? 3. What are the most effective methods for evaluating the generalization capabilities of NLP models? 4. While Large Language Models (LLMs) significantly enhance the generalization of NLP models, what are the key limitations of LLMs in this regard?

We received 128 submissions to the theme track during the review phase. Among these, 41 papers were accepted to the main conference and a further 33 to Findings of ACL 2025. The conference will also feature a panel discussion on the theme of Generalization, with the participation of leading experts in this area.

Best Paper Selection

ACL 2025 implemented the updated ACL award policy that seeks to expand the pool of work recognized as outstanding. In total 117 papers were nominated by the reviewers, area chairs and senior area chairs for best paper consideration. The best paper committee assessed these papers to select the best papers (featuring $\leq 0.6\%$ of accepted papers), outstanding papers (featuring $\leq 2.5\%$ of accepted papers), and special awards for social impact and best resource. Based on the review by the Best Paper committee, 39 papers have been selected for awards in the above categories. Separately, the senior area chairs also nominated their favourite papers as SAC Highlights.

In addition, we have awards for test of time award for a paper published in TACL in 2013 or 2014 and the best paper award for a paper published in TACL in 2024. The final selection was made by the best paper committee, and the winners will be announced during the closing ceremony. The ACL 2025 Best Papers will also be given an opportunity to present their work in the closing ceremony.

Program Composition & Presentation Modes

Based on feedback from the conference support staff and the Underline team after ACL 2025, we decided to hold the virtual presentations sessions during the main conference. This enables us to align the virtual sessions with time slots when in-person participants are available. This approach allows virtual attendees to participate concurrently with the physical event, avoiding the need for organizers and attendees to engage with the conference twice and separately. This year, 218 main conference papers were selected for oral presentations by the program chairs, with the goal of creating a well-rounded program featuring a diverse set of topics instead of selecting papers based on their review scores. This year we have introduced a panel section of the conference. Out of the selected oral presenters, 25 will have an opportunity to not only present their work but also participate in a panel discussion. We have lined up five panels on the theme of: Generalisation in NLP, LLM alignment, Human-centred NLP, Interpretability and model analysis, Multilinguality and language diversity. In addition to the main conference papers, the ACL program also includes 18 papers accepted by Computational Linguistics and 42 papers accepted by Transactions of the ACL (TACL). Among these, 11 journal papers will be presented in-person as oral presentations, thematically distributed across appropriate sessions.

This year the conference will also for the first time feature a dedicated Findings poster session with reception. All Findings presenters have been assigned poster presentations in this session or alongside other posters for main conference papers in the same track. Rounding out the program are dedicated sessions for the demonstrations track and the student research workshop.

Keynotes and Panel

This year’s program features an impressive lineup of two keynote presentations:

- Prof. Luke Zettlemoyer from Paul G. Allen School of Computer Science Engineering at the University of Washington, and a Senior Research Director at Meta, will share his insights on “Rethinking Pretraining: Data and Architecture.”
- Prof. Verena Rieser from Google DeepMind will present on “Whose Gold? Re-imagining Alignment for Truly Beneficial AI.”

Alongside these keynotes, we are thrilled to host a panel discussion aiming to answer the question Can large language models (LLMs) generalize?. Our esteemed panelists include:

- Prof. Eduard Hovy, University of Melbourne (who will also act as the panel chair).
- Prof. Mirella Lapata, University of Edinburgh
- Prof. Yue Zhang, Westlake University
- Prof. Dan Roth, University of Pennsylvania and Oracle

This diverse group of panelists will provide a comprehensive view of the latest trends and challenges in Generalisation of NLP in the Large Language Models era.

We hope you enjoy this year's diverse and engaging program!

Ekaterina Shutova (University of Amsterdam and Stanford University)

Mohammad Taher Pilehvar (Cardiff University and Tehran Institute for Advanced Studies)

Joyce Nakatumba-Nabende (Makerere University)

Wanxiang Che (Harbin Institute of Technology)

ACL 2025 Program Co-Chairs

ACL Business and Event Manager Message

On behalf of the ACL Executive team, I would like to extend my sincere gratitude to everyone who contributed to making ACL 2025 a reality.

A special thank you to our Local Committee—**Benjamin Roth** and **Dagmar Gromann**—as well as **Markus Heigl** from the Vienna Convention Bureau, our destination partner Meeting Destination Vienna, for their exceptional work throughout the submission and selection process. Your detailed planning and unwavering support helped bring ACL to this beautiful and historic city.

I also wish to thank the **ACL Boards** for their thoughtful review of venue bids and their decision to select Vienna, Austria as the host city. A particular note of appreciation goes to **David Yarowsky** for his guidance and efforts in finalizing our venue contracts. To our organizational partners and extended team:

- **ACL Onsite Team**
Megan Haddad – Event Assistant & ACL Admin
- **Auxiliary Staff**
Lina Staggs, Sally Stevenson, Brandy Dorsey, Karen Maloy, Jason Staggs
- **Support Staff – Lee Hartman & Sons**
Jon Dorsey – AV Director
Trevor Laffoon – AV Tech
- **Underline – Hybrid Event Partner**
Damira Mrsic – Content & Hybrid Manager
- **On-Site Support Team**
Jernej Masnec, Borna Bevanda, Leticija Maričić, Petra Filipčić, Petra Vižintin, Rafael Grabovica, Lucijan Prpic, Nikola Pesut, Astrid Plišo, Kristijan Varga, Filip Cigic, Mihael Herenda, Petar Vodinac, Roko Mirković
- **Austria Center Vienna**
Friedrich Gehmacher – for your guidance and help in contracting the venue
Victoria Wozniak – for managing the on-site logistics with care and precision
- **Vendor Partners**
Motto Catering
Stand Out (Exhibition Services)
XEST(Print Services)
Marcelo Viridian (Graphics)
- **Hotel Partners Meliá Vienna**
NH Danube City
Spark by Hilton Vienna Donaustadt

Finally, I want to express my heartfelt thanks to the entire **organizing and program committees**. Your dedication, long hours, and often sleepless nights have made ACL 2025 possible. Your commitment to academic excellence and community engagement is what makes this event a continued success year after year.

I look forward to a memorable and enriching conference here in Vienna, and I thank each of you for being a part of it. With gratitude,

Jennifer Rachford

Business and Event Manager

Organizing Committee

General Chair

Roberto Navigli, Sapienza University of Rome & Babelscape

Program Chairs

Wanxiang Che, Harbin Institute of Technology

Joyce Nakatumba-Nabende, Makerere University

Ekaterina Shutova, University of Amsterdam

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

Local Chairs

Benjamin Roth, University of Vienna

Dagmar Gromann, University of Vienna

Workshop Chairs

Terra Blevins, University of Vienna

Christophe Gravier, Université Jean Monnet

Tutorial Chairs

Yuki Arase, Tokyo Institute of Technology

David Jurgens, University of Michigan

Fei Xia, University of Washington

Demonstration Chairs

Pushkar Mishra, Google DeepMind / Cambridge

Smaranda Muresan, Barnard College, Columbia University

Tao Yu, The University of Hong Kong

Student Research Workshop Chairs

Zhu Liu, Tsinghua University

Mingyang Wang, LMU

Jin Zhao, Brandeis University

Student Research Workshop Chair: Faculty Advisors

Lea Frermann, University of Melbourne

Daniel Hershcovich, University of Copenhagen

Tristan Miller, University of Manitoba

Publication Chairs

Pierpaolo Basile, University of Bari

Libo Qin, Central South University

Zhenghao Liu, Northeastern University

Handbook Chairs

Els Lefever, Ghent University
Qiongkai Xu, Macquarie University

Sponsorship Chairs

Raffaella Bernardi, University of Trento
Thomas Scialom, Meta AI

Diversity and Inclusion Chairs

Senja Pollak, Jožef Stefan Institute
Maria Ryskina, MIT
Shane Storks, University of Michigan
Hwaran Lee, NAVER

Publicity and Social Media Chairs

Anette Frank, Heidelberg University
Shruti Rijhwani, Google DeepMind
Horacio Saggion, Pompeu Fabra University

Website and Conference App Chairs

Xudong Han, LibrAI / MBZUAI
Alessandro Raganato, University of Milano-Bicocca

Ethics Chairs

Karën Fort, Université de Lorraine / LORIA
Bjorn Ross, University of Edinburgh

Student Volunteer Chairs

Pedro Henrique Luz de Araujo, University of Vienna
Eleonora Mancini, University of Bologna

Technical Open Review Chairs

Niket Tandon, Allen AI
Lizhen Qu, Monash University

Virtual Infrastructure Chairs

Manling Li, Northwestern University
Yang Liu, Microsoft
Avi Sil, IBM Research

Internal Communication Chairs

Sara Tonelli, FBK

Yiquan Wu, Zhejiang University

Industry Track Chairs

Yunhao Li, Adobe
Georg Rehm, DFKI

Best Paper Committee Chairs

Rada Mihalcea, University of Michigan
Roi Reichart, Technion - Israel Institute of Technology

Visa Chairs

Rexhina Biloshmi, Amazon AGI
Eleni Ilkou, TIB

Documentation Chair

Chenghua Lin, University of Manchester

ARR Editors in Chief

Jun Suzuki, Tohoku University
Jing Jiang, Australian National University
Xiaodan Zhu, Queen's University

Program Committee

Area Chairs

David Ifeoluwa Adelani, McGill University
Alan Akbik, Humboldt Universität Berlin
Raquel G. Alhama, University of Amsterdam, University of Amsterdam
Afra Alishahi, Tilburg University
Jacob Andreas, Massachusetts Institute of Technology
Maria Antoniak
Marianna Apidianaki, University of Pennsylvania, University of Pennsylvania
Jun Araki, Bosch Research
Ehsaneddin Asgari, Qatar Computing Research Institute and University of California, Berkeley
Edoardo Barba, University of Roma "La Sapienza"
Jasmijn Bastings, Google DeepMind
Beata Beigman Klebanov, Educational Testing Service
Ahmad Beirami
Yonatan Belinkov, Technion, Technion
Meriem Beloucif, Uppsala University
Jonathan Berant, Google and Tel Aviv University
Steven Bethard, University of Arizona
Alexandra Birch, University of Edinburgh
Jelke Bloem, University of Amsterdam
Danushka Bollegala, Amazon and University of Liverpool
Houda Bouamor, Carnegie Mellon University
Jordan Lee Boyd-Graber, University of Maryland, College Park
Emanuele Bugliarello, Google DeepMind
Jan Buys, University of Cape Town
Aoife Cahill, Datamir
Jose Camacho-Collados, Cardiff University
Hanjie Chen, Rice University
Mingda Chen, Meta FAIR
Danqi Chen, Princeton University
Christos Christodoulopoulos, Amazon
Arman Cohan, Yale University and Allen Institute for Artificial Intelligence
Simone Conia, Sapienza University of Rome
Ryan Cotterell, Swiss Federal Institute of Technology
Yiming Cui, iFLYTEK Research
Raj Dabre, Department of Computer Science, Indian Institute of Technology, Madras, Indian Institute of Technology, Madras and National Institute of Information and Communications Technology (NICT), National Institute of Advanced Industrial Science and Technology
Tejaswini Deoskar, Utrecht University
Leon Derczynski, NVIDIA and IT University of Copenhagen
Yuxiao Dong, Tsinghua University
Li Dong, Microsoft Research
Yue Dong, University of California, Riverside and McGill University
A. Seza Doğruöz, Ghent University
Greg Durrett, University of Texas at Austin
Ondrej Dusek, Charles University, Prague
Nouha Dziri
Desmond Elliott, Copenhagen University and University of Copenhagen
Guy Emerson, University of Cambridge
Alexander Fabbri, Scale AI

Anna Feldman, Montclair State University and Montclair State University
 Fuli Feng, University of Science and Technology of China
 Xiaocheng Feng
 Anjalie Field, Johns Hopkins University
 Michael Flor, Educational Testing Service
 Antske Fokkens, VU University Amsterdam
 Zhe Gan, Apple
 Albert Gatt, Utrecht University
 Mor Geva, Tel Aviv University and Google Research
 Marjan Ghazvininejad, Facebook
 Goran Glavaš, Julius-Maximilians-Universität Würzburg
 Yvette Graham
 Jiafeng Guo, Institute of Computing Technolgy, Chinese Academy of Sciences
 Xianpei Han, Institute of Software, CAS
 Peter Hase, Schmidt Sciences
 Devamanyu Hazarika, Facebook
 Yulan He, King's College London, University of London
 Mohammad Javad Hosseini, Google
 Neil Houlsby, Google
 Dieuwke Hupkes, Facebook
 Ignacio Iacobacci, Elm Europe
 Mohit Iyyer, University of Maryland, College Park
 Zhijing Jin, Department of Computer Science, University of Toronto
 Shafiq Joty, Nanyang Technological University and Salesforce.com
 Rabeeh Karimi Mahabadi, NVIDIA
 Frank Keller, University of Edinburgh
 Daniel Khashabi, Johns Hopkins University
 Najoung Kim, Boston University
 Svetlana Kiritchenko, National Research Council Canada
 Ekaterina Kochmar, Mohamed bin Zayed University of Artificial Intelligence
 Julia Kreutzer, Cohere for AI
 Kalpesh Krishna, Google DeepMind
 Ranjay Krishna, University of Washington
 Matthieu Labeau, Télécom ParisTech
 Anne Lauscher, Universität Hamburg
 Carolin Lawrence, NEC Laboratories Europe and NEC Laboratories Europe
 Wenqiang Lei, Sichuan University
 Martha Lewis, University of Amsterdam
 Junyi Jessy Li, University of Texas at Austin
 Zhenghua Li, Soochow University
 Dongxu Li, Salesforce.com
 Piji Li, Nanjing University of Aeronautics and Astronautics
 Bill Yuchen Lin, xAI and University of Washington
 Yiqun Liu, Tsinghua University
 Fangyu Liu, Google DeepMind
 Zhiyuan Liu, Tsinghua University
 Kang Liu, Institute of automation, Chinese academy of science, Chinese Academy of Sciences
 Edison Marrese-Taylor, The University of Tokyo and AIST, National Institute of Advanced Industrial Science and Technology
 Joshua Maynez, Google
 Dipendra Misra, Mosaic Research, Databricks
 Saif M. Mohammad
 Nafise Sadat Moosavi, University of Sheffield

Aaron Mueller, Northeastern University and Technion - Israel Institute of Technology
Philippe Muller, IRIT, University of Toulouse
Preslav Nakov, Mohamed bin Zayed University of Artificial Intelligence
Vivi Nastase, University of Geneva
Wolfgang Nejdl, L3S Research Center and Universität Hannover
Xuefei Ning, Tsinghua University
Debora Nozza, Bocconi University
Giorgio Di Nunzio, University of Padua
Nedjma Ousidhoum, Cardiff University
Alexandros Papangelis, Apple
Gabiella Pasi, University of Milano-Bicocca
Rebecca J. Passonneau, Pennsylvania State University
Viviana Patti, University of Turin
Hao Peng, University of Illinois Urbana-Champaign
Carla Perez-Almendros, Cardiff University
Tiago Pimentel, Department of Computer Science, ETHZ - ETH Zurich
Edoardo Ponti, University of Edinburgh
Simone Paolo Ponzetto, Universität Mannheim
Daniel Preotiuc-Pietro, Bloomberg
Emily Prud'hommeaux, Boston College
Peng Qi, Orby AI
Marek Rei, Imperial College London
Verena Rieser, Google DeepMind
Roland Roller, German Research Center for AI
Alla Rozovskaya
Michael Sejr Schlichtkrull, Queen Mary, University of London
Steven Schockaert, Cardiff University
Rico Sennrich, University of Zurich
Ehsan Shareghi, Monash University
Qinlan Shen, Oracle
Kevin Small, Amazon
Mihai Surdeanu, University of Arizona
Anders Søgaard, Copenhagen University
Chenhao Tan, University of Chicago
Hao Tan, Adobe Systems
Harish Tayyar Madabushi, University of Bath
Alberto Testoni, Amsterdam UMC
Trang Tran, N/A
Marcos Vinicius Treviso
Yulia Tsvetkov, Department of Computer Science, University of Washington
Daniel Varab, German Research Center for AI
Suzan Verberne, Universiteit Leiden
Karin Verspoor, Royal Melbourne Institute of Technology
Ivan Vulić, Google DeepMind and University of Cambridge
Henning Wachsmuth, Leibniz Universität Hannover
Byron C Wallace, Northeastern University, Brown University and Northeastern University
Qifan Wang, Meta AI
Tong Xiao, Northeastern University
Deyi Xiong, Tianjin University
Ruifeng Xu, Harbin Institute of Technology
Yadollah Yaghoobzadeh
Diyi Yang, Stanford University
Min Yang, Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences, Chinese

Academy of Sciences

Helen Yannakoudakis, King's College London, University of London

Huaxiu Yao, Department of Computer Science, University of North Carolina at Chapel Hill

Andrew Yates, Johns Hopkins University

Kai Yu, Shanghai Jiao Tong University

Luke Zettlemoyer, University of Washington, Facebook and Meta

Weinan Zhang, Harbin Institute of Technology

Meishan Zhang, Harbin Institute of Technology (Shenzhen), China

Yue Zhang, Westlake University

Qi Zhang, Fudan University

Xin Zhao, Renmin University of China

Yi Zhou, Cardiff University

Arkaitz Zubiaga, Queen Mary University of London

Anti-harassment Policy

ACL 2025 adheres to the ACL Anti-Harassment Policy. Any participant who experiences harassment or hostile behavior may contact any current member of the ACL Professional Conduct Committee or Jennifer Rachford, who is usually available at the registration desk of the conference. Please be assured that if you approach us, your concerns will be kept in strict confidence, and we will consult with you on any actions taken. The open exchange of ideas, the freedom of thought and expression, and respectful scientific debate are central to the aims and goals of a ACL conference. These require a community and an environment that recognizes the inherent worth of every person and group, that fosters dignity, understanding, and mutual respect, and that embraces diversity. For these reasons, ACL is dedicated to providing a harassment-free experience for participants at our events and in our programs. Harassment and hostile behavior are unwelcome at any ACL conference. This includes speech or behavior (including in public presentations and on-line discourse) that intimidates, creates discomfort, or interferes with a persons participation or opportunity for participation in the conference. We aim for ACL conferences to be an environment where harassment in any form does not happen, including but not limited to: harassment based on race, gender, religion, age, color, national origin, ancestry, disability, sexual orientation, or gender identity. Harassment includes degrading verbal comments, deliberate intimidation, stalking, harassing photography or recording, inappropriate physical contact, and unwelcome sexual attention. The ACL board members are listed at: <https://www.aclweb.org/portal/about>. The full policy and its implementation is defined at: https://aclweb.org/adminwiki/index.php/Anti-Harassment_Policy

Ethics Policy

ACL adopts the ACM Code of Ethics (<https://www.acm.org/code-of-ethics>) in the version adopted June 22nd, 2018, by the ACM Council. In its application to ACL, it is to be read in the contextually appropriate interpretation, e.g., ACM member is to be read as ACL member. Sec 4.2 should be read as follows: 4.2 Treat violations of the Code as inconsistent with membership in the ACL. Each ACL member should encourage and support adherence by all members of the CL/NLP community regardless of ACL membership. ACL members who recognize a breach of the Code should consider reporting the violation to the ACL, which may result in remedial action. The open exchange of ideas, the freedom of thought and expression, and respectful scientific debate are central to the aims and goals of the ACL. These require a community and an environment that recognizes the inherent worth of every person and group, that fosters

dignity, understanding, and mutual respect, and embraces diversity. For these reasons, ACL is dedicated to providing a harassment-free experience for all the members, as well as participants at our events and in our programs. Harassment and hostile behavior are unwelcome at any ACL conference, associated event, or in ACL-affiliated online discussions. This includes speech or behavior that intimidates, creates discomfort, or interferes with a person's participation or opportunity for participation in a conference or an event. We aim for ACL-related activities to be an environment where harassment in any form does not happen, including but not limited to: harassment based on race, gender, religion, age, color, appearance, national origin, ancestry, disability, sexual orientation, or gender identity. Harassment includes degrading verbal comments, deliberate intimidation, stalking, harassing photography or recording, inappropriate physical contact, and unwelcome sexual attention. The policy is not intended to inhibit challenging scientific debate, but rather to promote it by ensuring that all are welcome to participate in the shared spirit of scientific inquiry. Vexatious complaints and willful misuse of this procedure will render the complainant subject to the same sanctions as a violation of the anti-harassment policy. It is the responsibility of the community as a whole to promote an inclusive and positive environment for our scholarly activities. In addition, anyone who experiences harassment or hostile behavior may contact any member of the Professional Conduct Committee. Members of this committee are instructed to keep any such contact in strict confidence, and those who approach the committee will be consulted before any actions are taken.

Meal Info

Breakfast

July 27 - August 1

Breakfast is not provided, you can check with your local hotel and near by restaurants.

Breaks †

July 27 - August 1

Coffee, tea, pastry, and fruit are provided late morning and midafternoon.

Lunch

July 27 - August 1

Lunch is not provided, however, the Austria Center Vienna will have lunch options available for purchase. See detailed agenda for times.

Dinner

July 27 - August 1

Dinner is not provided, but there are plenty of cafes, and restaurants within the area.

Social Events ††

Welcome Reception, July 27, 18:30 - 20:30

Light canapes, and a drink ticket will be provided on Sunday, July 27, 2025, at the Welcome Reception. It will be held on Level 0 Hall E of the Austria Center Vienna.

Findings Reception, Monday, July 28, 2025 †††

Light canapes, and refreshments will be provided on Monday, July 28, 2025, at the Findings Reception. It will be held on Level 0 Hall E of the Austria Center Vienna.

Social Event, July 29, 19:00 - 22:00

Social Event Gala Dinner International Buffet Dinner and a drink ticket will be provided on Tuesday, July 29, 2025, at the Social Event Gala Dinner (Social Gala). Located in Level 2 Hall A of the Austria Center Vienna.

†: Included in Main Conference, Tutorial, and Workshop registrations.

††: Included in the Main Conference registration only, and can be added on for Guests, Tutorials, Workshops, and Exhibitors to attend at the Registration Solutions Desk or through your YesEvents registration login. **No admission without an entry ticket**

†††: Included in the Main Conference registration only. No guest add ons for this reception

Social Events

ACL 2025 Welcome Reception

Venue: Austria Center Vienna, Level 0 Hall E

Time: Sunday, July 27, 2025 18:30 - 20:30

Kick off **ACL 2025** with an evening of networking, refreshments, and engaging conversations at our **Welcome Reception!** This gathering is the perfect opportunity to connect with fellow attendees, reconnect with colleagues, and immerse yourself in the vibrant natural language processing community before the conference begins in full.

Enjoy a relaxed atmosphere with light hors d'oeuvres and drinks as you meet researchers, practitioners, and industry leaders from around the world. Whether you're a first-time attendee or a returning participant, this reception sets the stage for an inspiring and productive conference experience.

Important Information:

Attendance is included for all Main Conference registrants. If you are not registered for the Main Conference but would like to attend, you may add this event to your registration at the Registration Solution Desk located on the lobby level.

We look forward to welcoming you to **ACL 2025!**

Warm regards,

The ACL 2025 Organizing Committee

ACL 2025 Social Event

Venue: Austria Center Vienna, Level 2 Hall A

Time: Tuesday, July 29, 2025 19:00 - 22:00

Join us for an unforgettable evening celebrating the spirit of Vienna, Austria. A perfect fusion of refinement, cultural exploration, and high-spirited fun, this event promises unforgettable moments for all in attendance.

Entertainment Lineup: Enjoy the graceful artistry of Waltz Ballet An exciting introduction to Voigas-plattler, an energetic folk dance that showcases rhythmic precision and vibrant footwork. An electrifying performance by the Sugar Swing Cats.

Enjoy a delicious buffet dinner featuring regional flavors. **Beer & wine tickets** available for guests. This is your chance to unwind, meet fellow attendees, and experience the best of Vienna's artistic charm. We look forward to celebrating with you!

Important Information: **Attendance is included** for all Main Conference registrants. If you are not registered for the Main Conference but would like to attend, you may **add this event to your registration** at the Registration Solutions Desk.

Warm regards,
The ACL 2025 Organizing Committee

Keynotes

Monday, July 28

Rethinking Pretraining: Data and Architecture**Luke Zettlemoyer**

Monday, July 28, 09:30 - 10:30

Abstract: Large language model training follows a standard pipeline: tokenization, pretraining, possibly mid-training, and post training or alignment. Despite its wild success, we understand relatively little about this recipe and are almost certainly missing many opportunities to improve it. In this talk, I will focus on three such cases. I'll describe our work on data efficient post training (e.g. LIMA, ALMA, and s1) where we argue that nearly all advanced model capabilities ultimately come from the pretraining data, even if effective alignment is still essential for controlling model behavior. I will also describe new methods for extracting more signal from the pretraining data, including new hierarchical architectures for byte-level language models (e.g. BLT) that are both tokenizer-free and scale better than traditional BPE-based methods, especially in the long tail. Finally, I will discuss decentralized, modular training algorithms

(e.g. BTM) that better isolate and control the influence of specific data on specific model components and behaviors. Together, these methods promise to simplify training and improve scaling, by centering and amplifying the influence of data in architecture design.

Bio: Luke Zettlemoyer is a Professor in the Paul G. Allen School of Computer Science & Engineering at the University of Washington, and a Senior Research Director at Meta. His research focuses on empirical methods for natural language semantics, and involves designing machine learning algorithms, introducing new tasks and datasets, and, most recently, studying how to best develop new architectures and self-supervision signals for pre-training. His honors include being elected ACL President, named an ACL Fellow, winning a PECASE award, an Allen Distinguished Investigator award, and multiple best paper awards. Luke was an undergrad at NC State, received his PhD from MIT and was a postdoc at the University of Edinburgh.

Tuesday, July 29

Whose Gold? Re-imagining Alignment for Truly Beneficial AI



Verena Rieser

Tuesday, July 29, 09:00 - 10:00

Abstract: Human feedback is often the "gold standard" for AI alignment, but what if this "gold" reflects diverse, even contradictory human values? This keynote explores the technical and ethical challenges of building beneficial AI when values conflict – not just between individuals, but also within them. My talk advocates for a dual expansion of the AI alignment framework: moving beyond a single, monolithic viewpoint to a plurality of perspectives, and transcending narrow safety and engagement metrics to promote comprehensive human well-being.

Bio: Verena Rieser is a Senior Staff Research Scientist at Google DeepMind, where she founded the VOICES team (Voices-of-all in alignment). Her team is a core contributor to Gemini with a mission to enhance model safety and usability for diverse communities. Verena has pioneered work in data-driven multimodal Dialogue Systems and Natural Language Generation, encompassing conversational RL agents, faithful data-to-text generation, spoken language understanding, evaluation methodologies, and applications of AI for societal good. Verena previously directed the NLP lab as a full professor at Heriot-Watt University, Edinburgh, and held a Royal Society Leverhulme Senior Research Fellowship. She earned her PhD from Saarland University.

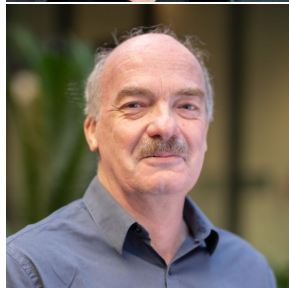
Invited Talks

6

Panel Discussion

Monday, July 28

Panel Discussion: Generalization of NLP Models



Panel: Mirella Lapata, Dan Roth, Yue Zhang

Moderator: Eduard Hovy

Monday, July 28, 16:30 - 17:30

Mirella Lapata is professor of natural language processing in the School of Informatics at the University of Edinburgh. Her research focuses on getting computers to understand, reason with, and generate natural language. She is the first recipient (2009) of the British Computer Society and Information Retrieval Specialist Group (BCS/IRSG) Karen Sparck Jones award and a Fellow of the Royal Society of Edinburgh, the ACL, and Academia Europaea. Mirella has also received best paper awards in leading NLP conferences and has served on the editorial boards of the Journal of Artificial Intelligence Research, the Transactions of the ACL, and Computational Linguistics. She was president of SIGDAT (the group that organizes EMNLP) in 2018. She has been awarded an ERC consolidator grant, a Royal Society Wolfson Research Merit Award, and a UKRI Turing AI World-Leading Researcher Fellowship.

Dan Roth is the Eduardo D. Glandt Distinguished Professor at the University of Pennsylvania and Chief AI Scientist at Oracle. Until June 2024 Dan was a VP/Distinguished Scientist at AWS AI where he led the scientific effort behind Amazon's first-generation GenAI products, including Titan Models, Amazon Q, and Amazon Bedrock. Dan is a Fellow of the AAAS, ACM, AAAI, and ACL, and a recipient of the IJCAI John McCarthy Award "for major conceptual and theoretical advances in the modeling of natural language understanding, machine learning, and reasoning." He has published broadly in natural language processing, machine learning, knowledge representation and reasoning, and learning theory, was the Editor-in-Chief of the Journal of Artificial Intelligence Research (JAIR) and has served as a Program Chair and Conference Chair for the major conferences in his research areas. Roth has been involved in several ML/NLP/GenAI startups in domains that range from legal and compliance to health care. Dan received his B.A Summa cum laude in Mathematics from the Technion, Israel and his Ph.D. in Computer Science from Harvard University in 1995.

Yue Zhang is a tenured Professor at Westlake University (<https://frcchang.github.io>). His research interests include fundamental NLP and its machine learning algorithms, and his recent research focuses on LLM reasoning and AI scientist. His major contributions to the field include machine learning algorithms for structured prediction (e.g., parsing and IE), neural NLP models (i.e., lattice and graph LSTM), and generalization for NLP/LM (e.g., OOD and logical reasoning). He co-authored the Cambridge University Press book "Natural Language Processing – a Machine Learning Perspective" and served as a PC co-chair for CCL 2020 and EMNLP 2022, test-of-time award committee co-chairs for ACL 2024 and 2025, action editor for TACL, and associate editor for TASLP, TALLIP, TBD, and CSL. He won the best paper awards of IALP 2017 and COLING 2018, best paper honorable mention of SemEval 2020, and best paper nomination for ACL 2018 and ACL 2023.

Eduard Hovy is the Executive Director of Melbourne Connect (a research and tech transfer centre at the University of Melbourne), a professor at the University of Melbourne's School of Computing and Information Systems, and an adjunct professor at the Language Technologies Institute in the School of Computer Science at Carnegie Mellon University. In 2020–21 he served as Program Manager in DARPA's Information Innovation Office (I2O), where he managed programs in Natural Language Technology and Data Analytics. Dr. Hovy completed a Ph.D. in Computer Science (Artificial Intelligence) at Yale University in 1987 and was awarded honorary doctorates from the National Distance Education University (UNED) in Madrid in 2013 and the University of Antwerp in 2015. He is one of the initial 17 Fellows of the Association for Computational Linguistics (ACL) and is also a Fellow of the Association for the Advancement of Artificial Intelligence (AAAI). Dr. Hovy's research focuses on computational semantics of language and addresses various areas in Natural Language Processing and Data Analytics, including in-depth machine reading of text, information extraction, automated text summarization, question answering, the semi-automated construction of large lexicons and ontologies, and machine translation. In early 2025 his Google h-index was 109, with about 68,000 citations. Dr. Hovy is the author or co-editor of eight books and over 400 technical articles and is a popular invited speaker. He regularly co-taught Ph.D.-level courses and has served on Advisory and Review Boards for both research institutes and funding organizations in Germany, Italy, Netherlands, Ireland, Singapore, and the USA. From 2003 to 2015 he was co-Director of Research for the Department of Homeland Security's Center of Excellence for Command, Control, and Interoperability Data Analytics, a distributed cooperation of 17 universities. In 2001 Dr. Hovy served as

President of the international Association of Computational Linguistics (ACL), in 2001–03 as President of the International Association of Machine Translation (IAMT), and in 2010–11 as President of the Digital Government Society (DGS).

Tutorials

Overview

Sunday, July 27, 2025

07:30 - 20:30	Registration	
09:00 - 12:30	Morning Tutorials	
	<i>Tutorial 1: Inverse Reinforcement Learning Meets Large Language Model Alignment</i>	<i>Level 2 Hall B</i>
	Mihaela van der Schaar and Hao Sun	
	<i>Tutorial 2: Eyetracking and NLP</i>	<i>Level 1 Hall M</i>
	David Reich, Omer Shubi, Lena Jäger and Yevgeni Berzak	
	<i>Tutorial 3: Uncertainty Quantification for Large Language Models</i>	<i>Level 2 Hall C</i>
	Artem Shelmanov, Maxim Panov, Ekaterina Sergeevna Fadeeva, Artem Vazhentsev, Roman Konstantinovich Vashurin and Timothy Baldwin	
	<i>Tutorial 4: Human-AI Collaboration: How AIs Augment Human Teammates</i>	<i>Level 1 Hall N</i>
	Tongshuang Wu, Diyi Yang, Kyle Lo and Marti A. Hearst	
14:00 - 17:30	Afternoon Tutorials	
	<i>Tutorial 5: Navigating Ethical Challenges in NLP: Hands-on strategies for students and researchers</i>	<i>Level 1 Hall M</i>
	Luciana Benotti, Fanny DuceL, Karen Fort, Guido Ivetta, Zhijing Jin, Min-Yen Kan, Seunghun Lee, Margot Mieskes, Minzhi Li, and Adriana Pagano	
	<i>Tutorial 6: NLP for Counterspeech against Hate and Misinformation</i>	<i>Level 1 Hall N</i>
	Daniel Russo, Helena Bonaldi, Yi-Ling Chung, Gavin Abercrombie and Marco Guerini	

Tutorial 7: Synthetic Data in the Era of Large Language Models *Level 2 Hall B*
Vijay Viswanathan, Xiang Yue, Alisa Liu, Yizhong Wang and Graham Neubig

Tutorial 8: Guardrails and Security for LLMs: Safe, Secure, and Controllable Steering of LLM Applications *Level 2 Hall C*
Traian Rebedea, Leon Derczynski, Shaona Ghosh, Makesh Narasimhan Sreedhar, Faeze Brahman, Liwei Jiang, Bo Li, Yulia Tsvetkov, Christopher Parisien and Yejin Choi

18:30 - 20:30

Welcome reception

Tutorial 1

T1: Inverse Reinforcement Learning Meets Large Language Model Alignment

Mihaela van der Schaar and Hao Sun

Large Language Model (LLM) alignment remains one of the most critical challenges in reinforcement learning. As the success of models like DeepSeek-R1 demonstrates, improving alignment requires better architectures and a deeper understanding of reinforcement learning (RL) and reward modeling. This tutorial explores the connection between Inverse Reinforcement Learning (IRL) and LLM alignment, offering a structured roadmap for researchers and practitioners.

We frame LLM alignment as an inverse RL problem, contrasting traditional reinforcement learning with inverse methods that infer rewards from human data. A key focus is on reward models, examining how they are constructed from various data sources, including mathematical reasoning, binary feedback, preference data, and demonstrations.

Beyond theory, we delve into infrastructure and practical implementation, showcasing how to efficiently evaluate IRL-based LLM alignment ideas in minutes. We conclude with insights from sparse-reward RL, covering reward shaping, credit assignment, and lessons from self-play.

By the end of this tutorial, attendees will gain a practical and theoretical understanding of LLM alignment through inverse RL, equipping them with the tools to build better-aligned models efficiently.

Tutorial 2

T2: Eyetracking and NLP

David Reich, Omer Shubi, Lena Jäger and Yevgeni Berzak

This is a tutorial on the growing research area that combines eyetracking during reading with NLP. The tutorial will outline how eye movements in reading can be leveraged for NLP, and, vice versa, how NLP methods can advance psycholinguistic modeling of eye movements in reading. We will cover four main themes: (i) fundamentals of eye movements in reading, (ii) experimental methodologies and available data, (iii) integrating eye movement data in NLP models, and (iv) using NLP for modeling eye movements in reading. The tutorial is tailored to NLP researchers and practitioners, and will provide attendees with the essential background for conducting research on the joint modeling of eye movements and text.

Tutorial 3

T3: Uncertainty Quantification for Large Language Models

Artem Shelmanov, Maxim Panov, Ekaterina Sergeevna Fadeeva, Artem Vazhentsev,
Roman Konstantinovich Vashurin and Timothy Baldwin

In this tutorial, we cover foundational concepts of uncertainty quantification (UQ) for LLMs, present cutting-edge techniques, demonstrate practical applications of UQ in various tasks, and equip researchers and practitioners with tools for developing new UQ methods and harnessing uncertainty in various contexts.

Tutorial 4

T4: Human-AI Collaboration: How AIs Augment Human Teammates

Tongshuang Wu, Diyi Yang, Kyle Lo and Marti A. Hearst

The continuous, rapid development of general-purpose models like LLMs suggests the theoretical possibility of AI performing any human task. Yet, despite the potential and promise, these models are far from perfect, excelling at certain tasks while struggling with others. The tension between what is possible and a model's limitations raises the general research question that has attracted attention from various disciplines: What is the best way to use AI to maximize its benefits? In this tutorial, we will review recent developments related to human-AI teaming and collaboration. To the best of our knowledge, our tutorial will be the first to provide a more integrated view from NLP, HCI, Computational Social Science, and Learning Science, etc., and highlight how different communities have identified the goals and societal impacts of such collaborations, both positive and negative. We will further discuss how to operationalize these Human-AI collaboration goals, and reflect on how state-of-the-art AI models should be evaluated and scaffolded to make them most useful in collaborative contexts.

Tutorial 5

T5: Navigating Ethical Challenges in NLP: Hands-on strategies for students and researchers

Luciana Benotti, Fanny Ducel, Karen Fort, Guido Ivetta, Zhijing Jin, Min-Yen Kan, Seunghun Lee, Margot Mieskes, Minzhi Li, and Adriana Pagano

With NLP research being rapidly productionized into real-world applications, it is important to be aware of and think through the consequences of our research. Such ethical considerations are important in both authoring and reviewing (e.g. privacy, consent, fairness, etc). This tutorial will equip participants with basic guidelines for thinking deeply about ethical issues and review common considerations that recur in NLP research. The methodology is interactive and participatory, including case studies and working in groups. Participants will gain practical experience on when to flag a paper for ethics review and how to write an ethical consideration section, which will be shared with the broader community. Importantly, the participants will be co-creating the tutorial outcomes and extending tutorial materials to share as public outcomes.

Tutorial 6

T6: NLP for Counterspeech against Hate and Misinformation

Daniel Russo, Helena Bonaldi, Yi-Ling Chung,
Gavin Abercrombie and Marco Guerini

- a. Introduction
 - b. Counterspeech in the real world: evidence from social science and the real world
 - c. NLP approaches to counterspeech against hate
 - d. NLP approaches to counterspeech for misinformation
 - e. Open challenges
 - f. Ethical considerations
 - g. Two guest talks from civil society/NGOs:
 - i. Cathy Buerger, Dangerous Speech Project ii. Simone Fontana, Facta News h. QA panel & Conclusion
-

Tutorial 7

T7: Synthetic Data in the Era of Large Language Models

Vijay Viswanathan, Xiang Yue, Alisa Liu, Yizhong Wang and Graham Neubig

This tutorial unifies the rapid progress in synthetic data generation, with a focus on language data. We will discuss fundamental algorithms and major applications for training large language models when manually-curated data is limited.

Tutorial 8

T8: Guardrails and Security for LLMs: Safe, Secure, and Controllable Steering of LLM Applications

Traian Rebedea, Leon Derczynski, Shaona Ghosh, Makes Narsimhan Sreedhar,
Faeze Brahman, Liwei Jiang, Bo Li, Yulia Tsvetkov,
Christopher Parisien and Yejin Choi

Pretrained generative models, especially large language models, provide novel ways for users to interact with computers. While generative NLP research and applications had previously aimed at very domain-specific or task-specific solutions, current LLMs and applications (e.g. dialogue systems, agents) are versatile across many tasks and domains. Despite being trained to be helpful and aligned with human preferences (e.g., harmlessness), enforcing robust guardrails on LLMs remains a challenge. And, even when protected against rudimentary attacks, just like other complex software, LLMs can be vulnerable to attacks using sophisticated adversarial inputs. This tutorial provides a comprehensive overview of key guardrail mechanisms developed for LLMs, along with evaluation methodologies and a detailed security assessment protocol - including auto red-teaming of LLM-powered applications. Our aim is to move beyond the discussion of single prompt attacks and evaluation frameworks towards addressing how guardrailing can be done in complex dialogue systems that employ LLMs.

We aim to provide a cutting-edge and complete overview of deployment risks associated with LLMs in production environments. While the main focus will be on how to effectively protect against safety and security threats, we also tackle the more recent topic of providing dialogue and topical rails, including respecting custom policies. We also examine the new attack vectors introduced by LLM-enabled dialogue systems, such as methods for circumventing dialogue steering.

Main Conference

Main Conference: Monday, July 28, 2025

08:00–18:30	Registration	<i>Level 0 Entrance Hall</i>
09:00–09:30	Session 1: Opening Session	<i>Level 2 Hall A–C</i>
09:30–10:30	Session 1: Keynote Luke Zettlemoyer	<i>Level 2 Hall A–C</i>
10:30–11:00	<i>Coffee Break</i>	<i>Foyers & Hall X5</i>
11:00–12:30	Session 2: Virtual Presentations	<i>Underline</i>
11:00–12:30	Session 2: In Person Posters (Main)	
	NLP Applications	<i>Level 0 Exhibit Halls X4 or X5</i>
	Resources and Evaluation	<i>Level 0 Exhibit Halls X4 or X5</i>
	Multimodality and Language Grounding	<i>Level 0 Exhibit Halls X4 or X5</i>
	Efficient Methods for NLP	<i>Level 0 Exhibit Halls X4 or X5</i>
	Generation	<i>Level 0 Exhibit Halls X4 or X5</i>
	Ethics, Bias and Fairness	<i>Level 0 Exhibit Halls X4 or X5</i>
	Information Retrieval and Text Mining	<i>Level 0 Exhibit Halls X4 or X5</i>
	Human-centered NLP	<i>Level 0 Exhibit Halls X4 or X5</i>
	Speech Processing	<i>Level 0 Exhibit Halls X4 or X5</i>
12:30–14:00	<i>Lunch Break</i>	
14:00–15:30	Session 3: In Person Orals	
	Machine Translation	<i>Level 1 1.85</i>
	Semantics	<i>Level 1 1.86</i>
	Ethics, Bias, and Fairness 1	<i>Level 1 M.2</i>
	Generation	<i>Level 1 M.1</i>
	Human-centered NLP	<i>Level 1 1.62</i>
	Information Retrieval and Text Mining	<i>Level 1 N.1</i>
	Linguistic Theories and Cognitive Modeling	<i>Level 1 1.61</i>
	Multimodality and Language Grounding 1	<i>Level 2 Hall C</i>
	NLP Applications 1	<i>Level 2 Hall B</i>
	Resources and Evaluation 1	<i>Level 2 Hall A</i>

	Theme track: Generalisation of NLP models 1	Level 1 1.15-16
	Student Research Workshop	Level 1 Hall N.2
	Industry	Level 1 Hall L.1 - L.3
	BoF	Level 1 1.31 - 1.32, 1.33
15:30–16:00	Coffee Break	Foyers & Hall X5
16:00–16:30	Session 4: CL Dissertation Award	Level 2 Hall A–C
16:30–17:30	Session 4: Panel	
	Topic: Generalization on NLP Models	Level 2 Hall A–C
18:00–19:30	Session 5: Virtual Presentations	Underline
18:00–19:30	Session 5: In Person Posters (Findings with Reception)	Level 0 Exhibit Halls X4 and X5

Main Conference: Tuesday, July 29, 2025

08:30–21:30	Registration	Level 0 Entrance Hall
09:00–10:00	Session 6: Keynote Verena Rieser	Level 2 Hall A–C
10:00–10:30	<i>Coffee Break</i>	Foyers & Hall X5
10:30–12:00	Session 7: Virtual Presentations	Underline
10:30–12:00	Session 7: In Person Posters (Main)	
	NLP Applications	Level 0 Exhibit Halls X4 or X5
	Resources and Evaluation	Level 0 Exhibit Halls X4 or X5
	Language Modeling	Level 0 Exhibit Halls X4 or X5
	Interpretability and Model Analysis	Level 0 Exhibit Halls X4 or X5
	Question Answering	Level 0 Exhibit Halls X4 or X5
	Machine Learning for NLP	Level 0 Exhibit Halls X4 or X5
	Multilingualism and Cross-Lingual NLP	Level 0 Exhibit Halls X4 or X5
	Semantics	Level 0 Exhibit Halls X4 or X5
	Student Research Workshop	Level 1 Hall N.2
	Industry	Level 1 Hall L.1–L.3
	BoF	Level 1 1.31–1.32, 1.33, 1.14
12:00–13:00	<i>Lunch Break</i>	
13:00–13:45	Session 8: Business Meeting	Level 2 Hall B
13:45–14:00	Break to Reset for orals	
14:00–15:30	Session 9: In Person Orals	
	Dialogue and Interactive Systems	Level 1 N.1
	Interpretability and Model Analysis 1	Level 1 M.2
	Language Modeling 1	Level 2 Hall C
	Summarization and Information Extraction	Level 1 N.2
	Efficient Methods for NLP 1	Level 1 M.1
	Resources and Evaluation 2	Level 2 Hall B
	Phonology, Morphology, and Syntax	Level 1 1.86
	Sentiment, Style and Argumentation Analysis	Level 1 1.62
	Multilingualism and Cross-Lingual NLP	Level 1 1.15–16
	Speech Processing	Level 1 1.61
	Computational Social Science and Cultural Analytics	Level 1 1.85
	Industry	Level 1 Hall L.1–L.3
	BoF	Level 1 1.31–1.32, 1.33, 1.14
15:30–16:00	<i>Coffee Break</i>	Foyers & Hall X5
16:00–17:30	Session 10: Virtual Presentations	Underline
16:00–17:30	Session 10: In Person Posters (Main)	
	NLP Applications	Level 0 Exhibit Halls X4 or X5
	Resources and Evaluation	Level 0 Exhibit Halls X4 or X5
	Multimodality and Language Grounding	Level 0 Exhibit Halls X4 or X5
	Efficient Methods for NLP	Level 0 Exhibit Halls X4 or X5
	Generation	Level 0 Exhibit Halls X4 or X5
	Ethics, Bias, and Fairness	Level 0 Exhibit Halls X4 or X5
	Dialogue and Interactive Systems	Level 0 Exhibit Halls X4 or X5
	Theme track: Generalisation of NLP models	Level 0 Exhibit Halls X4 or X5

	Machine Translation	<i>Level 0 Exhibit Halls X4 or X5</i>
	Summarization and Information Extraction	<i>Level 0 Exhibit Halls X4 or X5</i>
	BoF	<i>Level 1 1.31–1.32, 1.33, 1.14</i>
19:00–22:00	<i>Social Event for all Registered Full, Main, or Social Event add-ons</i>	<i>Level 2 Hall A</i>

Main Conference: Wednesday, July 30, 2025

08:30–17:30	Registration	<i>Level 0 Entrance Hall</i>
09:00–10:30	Session 11: In Person Orals	
	Panel 1: Generalisation of NLP models	<i>Level 2 Hall A</i>
	Panel 2: LLM alignment	<i>Level 2 Hall B</i>
	Panel 3: Human-centred NLP	<i>Level 2 Hall C</i>
	NLP Applications 2	<i>Level 1 M.1</i>
	Multimodality and Language Grounding 2	<i>Level 1 M.2</i>
	Panel 4: Interpretability and model analysis	<i>Level 1 N.1</i>
	Panel 5: Multilinguality and language diversity	<i>Level 1 N.2</i>
	Machine Learning for NLP	<i>Level 1 1.61</i>
	Efficient Methods for NLP 3	<i>Level 1 1.62</i>
	Interpretability and Model Analysis 2	<i>Level 1 1.85</i>
	Question Answering	<i>Level 1 1.86</i>
	Language Modeling 2	<i>Level 1 1.15–1.16</i>
	Industry	<i>Level 1 Hall L.1–L.3</i>
	BoF	<i>Level 1 1.31–1.32, 1.33, 1.14</i>
10:30–11:00	<i>Coffee Break</i>	<i>Foyers & Hall X5</i>
11:00–12:30	Session 12: Virtual Presentations	<i>Underline</i>
11:00–12:30	Session 12: In Person Posters (Main, IND, SRW)	
	NLP Applications	<i>Level 0 Exhibit Halls X4 or X5</i>
	Resources and Evaluation	<i>Level 0 Exhibit Halls X4 or X5</i>
	Language Modeling	<i>Level 0 Exhibit Halls X4 or X5</i>
	Interpretability and Model Analysis	<i>Level 0 Exhibit Halls X4 or X5</i>
	Question Answering	<i>Level 0 Exhibit Halls X4 or X5</i>
	Computational Social Science and Cultural Analytics	<i>Level 0 Exhibit Halls X4 or X5</i>
	Linguistic Theories and Cognitive Modeling	<i>Level 0 Exhibit Halls X4 or X5</i>
	Sentiment, Style and Argumentation Analysis	<i>Level 0 Exhibit Halls X4 or X5</i>
	Phonology, Morphology, and Syntax	<i>Level 0 Exhibit Halls X4 or X5</i>
	Discourse and Pragmatics	<i>Level 0 Exhibit Halls X4 or X5</i>
	Student Research Workshop: Closing and Best Papers	<i>Level 1 Hall N.2</i>
	Industry: Car-Panel / Awards Closing	<i>Level 1 Hall L.1–L.3</i>
	BoF	<i>Level 1 1.31–1.32, 1.33, 1.14</i>
12:30–14:00	<i>Lunch Break</i>	
14:00–14:30	Session 13: Presidential Address	<i>Level 2 Hall A–C</i>
14:30–14:50	Session 13: Test of Times	<i>Level 2 Hall A–C</i>
14:50–15:00	Session 13: ACL Distinguished Service Awards	<i>Level 2 Hall A–C</i>
15:00–16:00	Session 13: ACL Lifetime Achievement Award	<i>Level 2 Hall A–C</i>
16:00–16:30	<i>Coffee Break</i>	<i>Hall 3 Lobby Level</i>
16:30–17:30	Session 14: Best Paper Awards	<i>Level 2 Hall A–C</i>
17:30–18:00	Session 15: Closing Session	<i>Level 2 Hall A–C</i>

Workshops

Overview

During the days of the workshops, **Registration** will be held from 08:00- 16:30.

Workshop Posters will be on Level 1 Hall L1 - L3.

Thursday, July 31, 2025

Level 1 1.34	Workshop 1: 6th AfricaNLP Workshop: Multilingual and Multicultural-aware LLMs	p.49
Level 2 Hall B	Workshop 2: The 12th Workshop on Argument Mining	p.50
Level 2 2.17	Workshop 3: 2nd Workshop on Natural Language Processing meets Climate Change	p.51
Level 2 2.31	Workshop 4: Eighth Workshop on Fact Extraction and VERification (FEVER)	p.52
Level 2 Hall C	Workshop 5: GEM2 Workshop: Generation, Evaluation & Metrics	p.53
Level 1 1.15-16	Workshop 6: The 19th Linguistic Annotation Workshop (LAW XIX)	p.54
Level 1 Hall N1	Workshop 7: The 4th Workshop on NLP for Positive Impact	p.55
Level 1 1.61-62	Workshop 8: REALM: First Workshop for Research on Agent Language Models	p.56
Level 1 1.14	Workshop 9: SDP 2025: The 5th Workshop on Scholarly Document Processing	p.57
Level 1 1.31-32	Workshop 10: The Third Workshop on Social Influence in Conversations (SICon 2025)	p.58
Level 2 2.44	Workshop 11: Slavic NLP-2025: The 10th Biennial Workshop on Slavic NLP	p.59
Level 2 2.15	Workshop 12: The 4th Table Representation Learning Workshop at ACL 2025	p.60
Gather	Workshop 13: The Fourth Ukrainian Natural Language Processing Workshop (UNLP 2025)	p.61
Level 1 1.85-86	Workshop 14: The 20th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2025)	p.62
Level 1 Hall M1	Workshop 15: 29th Conference on Computational Natural Language Learning (CoNLL)	p.63
Level 1 Hall N2	Workshop 16: The 22nd International Conference on Spoken Language Translation (IWSLT 2025)	p.64
Level 1 Hall M2	Workshop 17: The 19th Workshop on Semantic Evaluation (Semeval-2025)	p.65

Friday, August 1, 2025

Gather	Workshop 13: The Fourth Ukrainian Natural Language Processing Workshop (UNLP 2025)	p.61
Level 1 1.85-86	Workshop 14: The 20th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2025)	p.62
Level 1 Hall M1	Workshop 15: 29th Conference on Computational Natural Language Learning (CoNLL)	p.63
Level 1 Hall N2	Workshop 16: The 22nd International Conference on Spoken Language Translation (IWSLT 2025)	p.64
Level 1 Hall M2	Workshop 17: The 19th Workshop on Semantic Evaluation (Semeval-2025)	p.65
Level 1 1.34	Workshop 18: The Second Workshop on Analogical Abstraction in Cognition, Perception, and Language (Analogy-ANGLE II)	p.66
Level 2 2.15	Workshop 19: BioNLP 2025 and Shared Tasks (BioNLP-ST 2025)	p.67
Level 2 Hall C	Workshop 20: The 6th Workshop on Gender Bias in Natural Language Processing (GeBNLP 2025)	p.68
Level 1 1.31-32	Workshop 21: The First Workshop on Large Language Model Memorization (L2M2)	p.69
Level 2 Hall B	Workshop 22: First Workshop on LLM Security (LLMSEC)	p.70
Level 2 2.44	Workshop 23: MAGMaR: Multimodal Augmented Generation via Multimodal Retrieval	p.71
Level 1 1.14	Workshop 24: FieldMatters–SIGTYP	p.72
Level 1 Hall N1	Workshop 25: The 3rd workshop on Towards Knowledgeable Foundation Models	p.73
Level 2 2.31	Workshop 26: WikiNLP: Advancing Natural Language Processing for Wikipedia	p.74
Level 2 Hall 2.17	Workshop 27: The 9th Workshop on Online Abuse and Harms (WOAH 2025)	p.75
Level 1 1.61-62	Workshop 28: The First Joint Workshop on Large Language Models and Structure Modeling	p.76

Workshop 01

W01: 6th AfricaNLP Workshop: Multilingual and Multicultural-aware LLMs

David Adelani, Constantine Lignos, Tajuddeen Gwadabe, Bonaventure F. P. Dosso, Israel Abebe Azime, Everlyn Asiko Chimoto, Clemencia Siro, Henok Biadgign Ademteu, Happy Buzaaba, Hatem Haddad, Menno Van Zaanen, Mmasibidi Setaka, Andiswa Bukula, Nomsa Skosana

Thursday, July 31
Level 1 Hall 1.34

Description: The AfricaNLP workshop brings together a diverse group of researchers to explore solutions, collaborations, and innovation around enhancing LLMs' capabilities in African languages and ensuring cultural awareness in their applications.

Workshop 02

W02: The 12th Workshop on Argument Mining

Elena Chistova, Philipp Cimiano, Shohreh Haddadan,
Gabiella Lapesa, Ramon Ruiz-Dolz

Thursday, July 31
Level 2 Hall B

Description: Argument Mining (also known as “argumentation mining”) is an emerging research area within computational linguistics that started with focusing on automatically identifying and classifying argument elements, covering several text genres such as legal documents, news articles, online debates, scholarly data, and many more. In recent years, the field (broadly Computational Argumentation) has grown to explore argument quality and synthesis on many levels. The field offers practical uses such as argument-focused search and debating technologies, e.g., IBM Project Debater. The growing interest in computational argumentation has led to several tutorials at major NLP conferences.

Besides providing a forum to discuss and exchange cutting edge research in this field, a secondary goal of this year’s edition will be to broaden the disciplinary scope of the workshop by inviting other disciplines (e.g., (computational) social and political science, psychology, humanities) as well as other subareas of NLP to actively participate in the workshop and further shaping the field of argument mining. In particular, we would like to create synergies between the fields of argument mining and natural language reasoning.

Workshop 03

W03: 2nd Workshop on Natural Language Processing meets Climate Change

Dominik Stammbach, Tobias Schimanski, Jingwei Ni, Alba Sun, Alok Singh, Christopher Manning, Gaku Morio, Kalyan Dutia, Peter Henderson, Saeid Vaghefi, Veruska Muccione

Thursday, July 31
Level 2 Hall 2.17

Description: This workshop is dedicated to discuss how Natural Language Processing can be incorporated in Climate Change science, and help mitigating or adapting to climate change.

Workshop 04

W04: Eighth Workshop on Fact Extraction and VERification (FEVER)

Andreas Vlachos, Mubashara Akhtar, Michael Schlichtkrull, Christos Christodoulopoulos, Yulong Chen, Zhenyun Deng, Marek Strong, Arpit Mittal, Oana Cocarascu, Chenxi Whitehouse, James Thorne, Zhijiang Guo, Rami Aly, Rui Cao

Thursday, July 31
Level 2 Hall 2.31

Description: The Fact Extraction and Verification Workshop brings together researchers working on this topic as well as related ones such as recognizing textual entailment, question answering and argumentation mining.

Workshop 05

W05: Workshop on Meaningful, Efficient, and Robust Evaluation of LLMs and GEM: Natural Language Generation, Evaluation, and Metrics

Sebastian Gehrmann, Gabriel Stanovsky, Enrico Santus, Itay Itzhak, João Sedoc, Kaustubh Dhole, Michal Shmueli Scheuer, Miruna Clinciu, Ofir Arviv, Rotem Dror, Simon Mille, Yotam Perlitz, Oyvind Tafjord

Thursday, July 31
Level 2 Hall C

Description: Evaluating large language models (LLMs) is challenging. Running LLMs over medium or large scale corpus can be prohibitively expensive; they are consistently shown to be highly sensitive to prompt phrasing, and it is hard to formulate metrics which differentiate and rank different LLMs in a meaningful way. Consequently, the validity of the results obtained over popular benchmarks such as HELM or MMLU, lead to brittle conclusions (Sclar et al., 2024, Mizrahi et al., 2024, Alzahrani et al., 2024). We believe that meaningful, efficient, and robust evaluation is one of the cornerstones of the scientific method, and that achieving it should be a community-wide goal.

In this workshop we seek innovative research relating to the evaluation of LLMs and language generation systems in general. This includes, but is not limited to, robust, reproducible and efficient evaluation metrics, as well as new approaches for collecting evaluation data which can help in better differentiating between different systems and understanding their current bottlenecks.

Workshop 06

W06: The 19th Linguistic Annotation Workshop (LAW XIX)

Siyao Peng, Ines Rehbein, Amir Zeldes

Thursday, July 31
Level 1 Hall 1.15-16

Description: The Linguistic Annotation Workshop (LAW 2025) targets innovative research on linguistic annotation, focusing on subjectivity and the annotation of concepts from social sciences.

Workshop 07

W07: The 4th Workshop on NLP for Positive Impact

Ruyuan Wan, Katherine Atwell, Laura Biester, Angana Borah, Daryna Dementieva,
Oana Ignat, Neema Kotonya, Ziyi Liu, Steven Wilson, Jieyu Zhao

Thursday, July 31
Level 1 Hall N1

Description: The 4th Workshop on NLP for Positive Impact aims to advance innovative NLP research for positive societal impact, emphasizing responsible methods and diverse applications.

Workshop 08

W08: REALM: First Workshop for Research on Agent Language Models

Ehsan Kamalloo, Nicolas Gontier, Xing Han Lu, Shikhar Murty, Alexandre Lacoste,
Nouha Dziri, Hanna Hajishirzi, Graham Neubig

Thursday, July 31
Level 1 Hall 1.61-62

Description: The REALM workshop aims to bring together researchers, practitioners, and thought leaders to discuss and align on the current landscape, key challenges, and future directions of LLM Agents.

Workshop 09

W09: SDP 2025: The 5th Workshop on Scholarly Document Processing

Tirthankar Ghosal, Philipp Mayr, Aakanksha Naik, Amanpreet Singh, Anita de
Waard, Dayne Freitag, Georg Rehm, Sonja Schimmler, Dan Li

Thursday, July 31
Level 1 Hall 1.14

Description: The goal of this workshop is to provide a venue for addressing these challenges, as well as a platform for tasks and resources supporting the processing of scientific documents. Our long-term objective is to establish scholarly and scientific texts as an essential domain for NLP research, to supplement current efforts on web text and news articles.

Workshop 10

W10: The Third Workshop on Social Influence in Conversations (SICon 2025)

James Hale, Deuksin Kwon, Kushal Chawla, Ritam Dutt, Muskan Garg, Liang Qiu,
Alexandros Papangelis, Gale Lucas, Zhou Yu, Daniel Hershcovich

Thursday, July 31
Level 1 Hall 1.31-32

Description: Social influence (SI) is the change in an individual's thoughts, feelings, attitudes, or behaviors from interacting with another individual or a group. For example, a buyer uses SI skills to negotiate trade-offs and build rapport with the seller. SI is ubiquitous in everyday life, and hence, realistic human-machine conversations must reflect these dynamics, making it essential to model and understand SI in dialogue research systematically. This would improve SI systems' ability to understand users' utterances, tailor communication strategies, personalize responses, and actively lead conversations. These challenges draw on perspectives not only from NLP and AI research but also from Game Theory, Affective Computing, Communication, and Social Psychology.

Workshop 11

W11: Slavic NLP-2025: The 10th Biennial Workshop on Slavic NLP

Jakub Piskorski, Michał Marcińczuk, Preslav Nakov, Nikola Ljubešić,
Pavel Přibáň, Roman Yangarber

Thursday, July 31
Level 2 Hall 2.44

Description: Slavic NLP 2025 is a Workshop that addresses Natural Language Processing (NLP) for the Slavic languages.

The goal of this Workshop is to bring together researchers from academia and industry working on NLP for Slavic languages. In particular, the Workshop will serve to stimulate research and foster the creation of tools and resources for these languages. We also organize the shared task on the Detection and Classification of Persuasion Techniques in Texts for Slavic Languages.

Workshop 12

W12: The 4th Table Representation Learning Workshop at ACL 2025

Madelon Hulsebos, Shuaichen Chang, Wenhui Chen, Filip Gralinski,
Qian Liu, Huan Sun

Thursday, July 31
Level 2 Hall 2.15

Description: "Tables are a promising modality for representation learning and generative models with too much application potential to ignore. The Table Representation Learning (TRL) workshop is the premier venue in this emerging research area and has three main goals: (1) Motivate structured data (e.g. tables) as a primary modality for representation and generative models and advance the area further. (2) Showcase impactful applications of pretrained table models and identify open challenges for future research, with a particular focus on progress in NLP for this edition at ACL in 2025. (3) Foster discussion and collaboration across the NLP, ML, IR and DB communities."

Workshop 13

W13: The Fourth Ukrainian Natural Language Processing Workshop (UNLP 2025)

Mariana Romanyshyn, Olena Nahorna, Oleksii Ignatenko, Andrii Hlybovets

Thursday, July 31 and Friday, August 1, 2025
Virtually on Gather

Description: The UNLP workshop is dedicated to the development of language resources, tools, and NLP solutions for the Ukrainian language. This event brings together professionals from academia and industry who work with Ukrainian or do cross-lingual research that can be applied to Ukrainian. Every year UNLP hosts a shared task and invited speakers who promote research on low-to-mid-resource languages.

Workshop 14

W14: The 20th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2025)

Ekaterina Kochmar, Marie Bexte, Andrea Horbach, Ronja Laarmann-Quante, Anaïs Tack, Bashar Alhafni, Victoria Yaneva, Jill Burstein, Zheng Yuan

Thursday, July 31 and Friday, August 1, 2025
Level 1 Hall 1.85-86

Description: The BEA Workshop is a leading venue for NLP innovation in the context of educational applications.

Workshop 15

W15: 29th Conference on Computational Natural Language Learning (CoNLL)

Gemma Boleda, Michael Roth, Emily Cheng, Selina Meyer, Snigdha Chaturvedi

Thursday, July 31 and Friday, August 1, 2025
Level 1 Hall M1

Description: CoNLL is a yearly conference organized by SIGNLL (ACL's Special Interest Group on Natural Language Learning), focusing on theoretically, cognitively, and scientifically motivated approaches to computational linguistics. In particular, it seeks to explore the interaction between theoretical issues in linguistics and cognition, on the one hand, and computational modeling, on the other.

Workshop 16

W16: The 22nd International Conference on Spoken Language Translation (IWSLT 2025)

Marcello Federico, Alex Waibel, Elizabeth Salesky, Jan Niehues, Sebastian Stüker,
Atul Kr. Ojha, Marine Carpuat

Thursday, July 31 and Friday, August 1, 2025
Level 1 Hall N2

Description: The International Conference on Spoken Language Translation (IWSLT) is an annual scientific conference dedicated to all aspects of spoken language translation. For more than 20 years, the conference has published research papers and organized key evaluation campaigns within the field, including the creation of key data suites, benchmarks, metrics and new tasks. IWSLT is the annual meeting of SIGSLT, the ACL-ISCA-ELRA Special Interest Group on Spoken Language Translation.

Workshop 17

W17: The 19th Workshop on Semantic Evaluation (Semeval-2025)

Sara Rosenthal, Aiala Rosá, Marcos Zampieri, Debanjan Ghosh

Thursday, July 31 and Friday, August 1, 2025
Level 1 Hall M2

Description: The Semantic Evaluation (SemEval) workshops focus on the evaluation and comparison of systems that analyze diverse semantic phenomena in text, with the aim of extending the current state of the art in semantic analysis and creating high quality annotated datasets in a range of increasingly challenging problems in natural language semantics.

Workshop 18

W18: The Second Workshop on Analogical Abstraction in Cognition, Perception, and Language (Analogy-ANGLE II)

Filip Ilievski, Giulia Rambelli, Marianna Bolognesi, Ute Schmid, Pia Sommerauer

Friday, August 1, 2025
Level 1 Hall 1.34

Description:

The Second Workshop on Analogical Abstraction in Cognition, Perception, and Language (Analogy-Angle II). Explore, model, and understand analogical reasoning in cognition, language, and computational models from an interdisciplinary perspective. Analogy-Angle II is a multidisciplinary workshop to advance research on analogical abstraction by bridging the fields of computational linguistics, artificial intelligence, and cognitive psychology. This workshop seeks to foster collaboration among researchers by providing a platform for sharing novel insights, benchmarks, methodologies, and analogy applications across disciplines.

Workshop 19

W19: BioNLP 2025 and Shared Tasks (BioNLP-ST 2025)

Dina Demner-Fushman, Sophia Ananiadou, Makoto Miwa, Jun-ichi Tsujii

Friday, August 1, 2025
Level 2 Hall 2.15

Description: BioNLP associated with ACL SIGBIOMED is an established primary venue for presenting research in language processing for the biomedical domains. The workshop has been running every year since 2002 and continues getting stronger. BioNLP truly encompasses the breadth of the domain and brings together researchers in biomedical and clinical NLP from all over the world.

Workshop 20

W20: The 6th Workshop on Gender Bias in Natural Language Processing (GeBNLP 2025)

Christine Basta, Marta R. Costa-jussà, Agnieszka Faléńska,
Debora Nozza, Karolina Stańczak

Friday, August 1, 2025
Level 2 Hall C

Description: The GeBNLP workshop provides a dedicated forum to address gender and other demographic biases in natural language processing models. It aims to foster awareness, share research on mitigating bias through data and algorithmic approaches, and build community consensus around standard evaluation tasks and metrics.

Workshop 21

W21: The First Workshop on Large Language Model Memorization (L2M2)

Robin Jia, Eric Wallace, Yangsibo Huang, Tiago Pimentel, Pratyush Maini, Verna Dankers, Johnny Wei, Pietro Lesci

Friday, August 1, 2025
Level 1 Hall 1.31-32

Description: The First Workshop on Large Language Model Memorization (L2M2), co-located with ACL 2025 in Vienna, seeks to provide a central venue for researchers studying LLM memorization from these different angles.

Workshop 22

W22: First Workshop on LLM Security (LLMSEC)

Leon Derczynski, Jekaterina Novikova, Muhao Chen

Friday, August 1, 2025
Level 2 Hall B

Description: Work on adversarially-induced failure modes of large language models, the conditions that lead to them, and their mitigations.

Workshop 23

W23: MAGMaR: Multimodal Augmented Generation via Multimodal Retrieval

Reno Kriz, Kenton Murray, Kate Sanders, Eugene Yang, Cameron Carpenter,
Francis Ferraro, Benjamin Van Durme

Friday, August 1, 2025
Level 2 Hall 2.44

Description: Vast amounts of information today is being stored as videos with minimal text metadata, necessitating further research around the efficient discovery, understanding, and synthesis of multimodal collections. To address this need, the workshop on Multimodal Augmented Generation via Multimodal Retrieval (MAGMaR) workshop will focus on two primary areas: (1) the retrieval of multimodal content, which spans text, images, audio, video, and multimodal data (e.g., image-language, video-language); and (2) retrieval-augmented generation, with an emphasis on multimodal retrieval and generation.

Workshop 24

W24: FieldMatters–SIGTYP: The 7th Workshop on Research in Computational Linguistic Typology and Multilingual NLP (SIGTYP) and the Fourth Workshop on NLP Applications to Field Linguistics (FieldMatters)

Michael Hahn, Ekaterina Vylomova

Friday, August 1, 2025
Level 1 Hall N1

Description: SIGTYP is the first dedicated venue for typology-related research and its integration in multilingual NLP.

Workshop 25

W25: The 3rd workshop on Towards Knowledgeable Foundation Models

Yuji Zhang, Xiaozhi Wang, Mor Geva, Chi Han, Shangbin Feng, Silin Gao, Sha Li, Manling Li, Heng Ji

Friday, August 1, 2025
Level 1 Hall 1.14

Description: Knowledge has been an important pre-requisite for a variety of AI applications, and is typically sourced from either structured knowledge sources such as knowledge bases and dictionaries or unstructured knowledge sources such as Wikipedia documents. More recently, researchers have discovered that language models already possess a significant amount of knowledge through pre-training: LLMs can be used to generate commonsense knowledge and factual knowledge context for question answering. While the results are encouraging, there are still lingering questions: Where does this knowledge come from? How much do language models know? Is this knowledge reliable? If some knowledge is wrong, can we fix it?

Workshop 26

W26: WikiNLP: Advancing Natural Language Processing for Wikipedia

Akhil Arora, Isaac Johnson, Lucie-Aimée Kaffee, Tzu-Sheng Kuo,
Tiziano Piccardi, Indira Sen

Friday, August 1, 2025
Level 2 Hall 2.31

Description: WikiNLP invites researchers to contribute novel uses of Wikimedia data or studies of the impact of Wikimedia data within the NLP community. We will discuss tensions around multilinguality and concerns raised by generative AI. We will also highlight successful approaches to developing tooling that can assist the Wikimedia community in maintaining and improving the breadth of the Wikimedia projects.

Workshop 27

W27: The 9th Workshop on Online Abuse and Harms (WOAH 2025)

Zeeraak Talat, Agostina Calabrese, Christine de Kock, Francielle Vargas, Flor
Miriam Plaza del Acro, Debora Nozza

Friday, August 1, 2025
Level 2 Hall 2.17

Description: WOAH, the 9th Workshop on Online Abuse and Harms, invites paper submissions from a diverse range of fields. Digital technologies have brought significant benefits to society, transforming how people connect, communicate, and interact. However, these same technologies have also enabled the widespread dissemination and amplification of abusive and harmful content, such as hate speech, harassment, and misinformation. Given the sheer volume of content shared online, addressing abuse and harm at scale requires the use of computational tools. Yet, detecting and moderating online abuse remains a complex task, fraught with technical, social, legal, and ethical challenges.

Workshop 28

W28: The First Joint Workshop on Large Language Models and Structure Modeling

Zeera Talat, Agostina Calabrese, Christine de Kock, Francielle Vargas, Flor
Miriam Plaza del Acro, Debora Nozza

Friday, August 1, 2025
Level 1 Hall 1.61-62

Description: The 1st Joint Workshop on Large Language Models and Structure Modeling (XLLM 2025) at ACL 2025 aims to encourage discussions and highlight methods for language structure modeling in the era of LLMs.

10

Poster Printing

XEST is the ACL 2025 preferred **poster printing** company. The daily **deadline** to order for next-day pick-up is 11 a.m.

Poster pick-up times will be daily, 08:00 - 15:00, Monday, July 28 - Friday, August 1 on Level 0, Hall 0.11



Last minute order? Scan the QR code below.
The daily deadline to order for next-day pick-up is 11 a.m.



Alternative near Austria Center Vienna: <https://www.iprinter.at/copyshop.html>

Alternative in the City Center: <https://diekopie.at/poster-plane-plakate/>

11

Local Guide

About Vienna

Architecturally, the city is still characterized primarily by the buildings around Vienna's Ringstrasse from the Wilhelminian period, but also by Baroque and Art Nouveau (Vienna Modernism and Vienna Secession). Through its role as the imperial capital and residence of the Austrian Empire from 1804 onwards, Vienna became a cultural and political center of Europe. The historic center of Vienna and Schönbrunn Palace are now UNESCO World Heritage sites. With around 7.5 million tourists annually and around 18.9 million overnight stays in 2024, Vienna is one of the most visited cities in Europe.

Language

German is considered the official language of Austria, however English is widely spoken throughout the country.

Safety

Vienna is generally considered a very safe place to visit, but like visiting any major city, be aware of your surroundings and of scams, know emergency numbers (listed below), keep digital copies of important documents (passport).

Power Plugs/Outlets

Outlets in Vienna are European Type F and the voltage is 230 volts, 50 Hz. This means you'll need a travel adapter or possible converter for devices from countries like the US, which have different plug types and voltages.

In case of an Emergency

European Emergency Number: 112 (Ambulance, Fire, Police)

Clothing

Plan on packing light, breathable summer clothing. Please note that in Vienna there can be very hot summer days and heavy rain. Water mists are provided at certain locations for cooling refreshment. The City of Vienna also provides free water from drinking fountains.

Currency & Living Costs

The currency used is the 'Euro' written as EUR or €.

Euro notes start from 5 € and go up to 200 €, whereas the coins have denominations of 1 cent, 2 cents, 5 cents, 10 cents, 20 cents, 50 cents, 1 €, and 2 €.

Average Cost of some items

- Business Lunch 25 €
- Cup of Coffee 4 €
- Filled Sandwich 5 €
- Mineral Water 2 €

Most stores accept credit and debit cards

Water

The tap water in Vienna has a very high quality and is safe to drink.

Vienna Sights and Tourism Info Desk

The Vienna Tourism Board will provide a tourism info desk directly at the conference venue from July 28 to 30, 2025 for your convenience. The provided services include:

- Tourist information about Vienna

- Event details and recommendations
- Assistance with hotel bookings
- Various brochures about Vienna, including city maps in multiple languages

Feel free to also explore all attractive sights and events in Vienna around the time of the conference here.

The official Vienna City Card provides you with discounts for public transportation and tourist attractions. Visit <https://www.viennacitycard.at/en/> for more information.

Popular Destinations

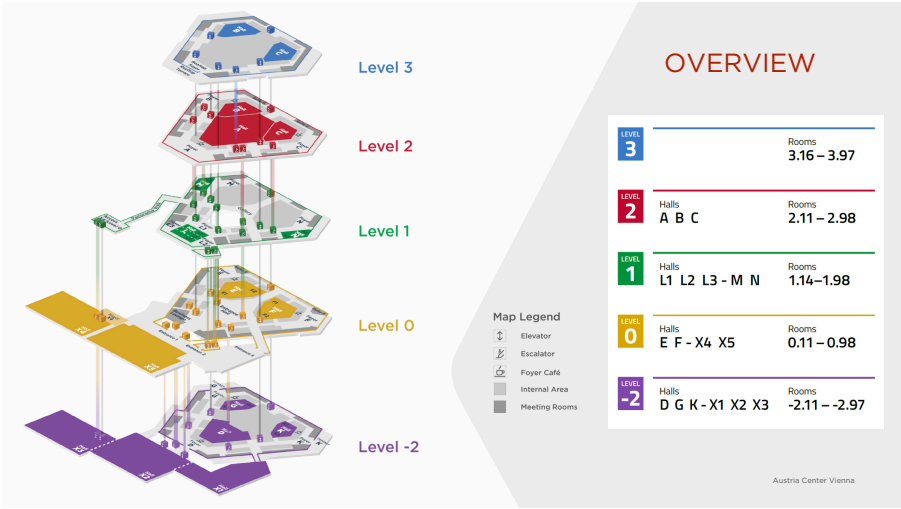
- Albertina (Art museum: Monet, Picasso and others, <https://www.albertina.at/en/home>)
- Kunsthistorisches Museum (Museum of Art History, very large collection: Tizian, Raffael, Caravaggio, Dürer, Bruegel, Rubens, Vermeer, Rembrandt, ...) <https://www.khm.at/>
- Riesenrad (Giant ferris wheel in the Prater amusement park) <https://wienerriesenrad.com/en/home-2>
- Stephansdom (Cathedral, <https://www.stephanskirche.at/visit.php>)
- Schloss Schönbrunn (Castle, <https://www.schoenbrunn.at/en/>)
- Schloss Belvedere (Castle and art museum, holds Klimt's famous work "The Kiss" <https://www.belvedere.at/en>)

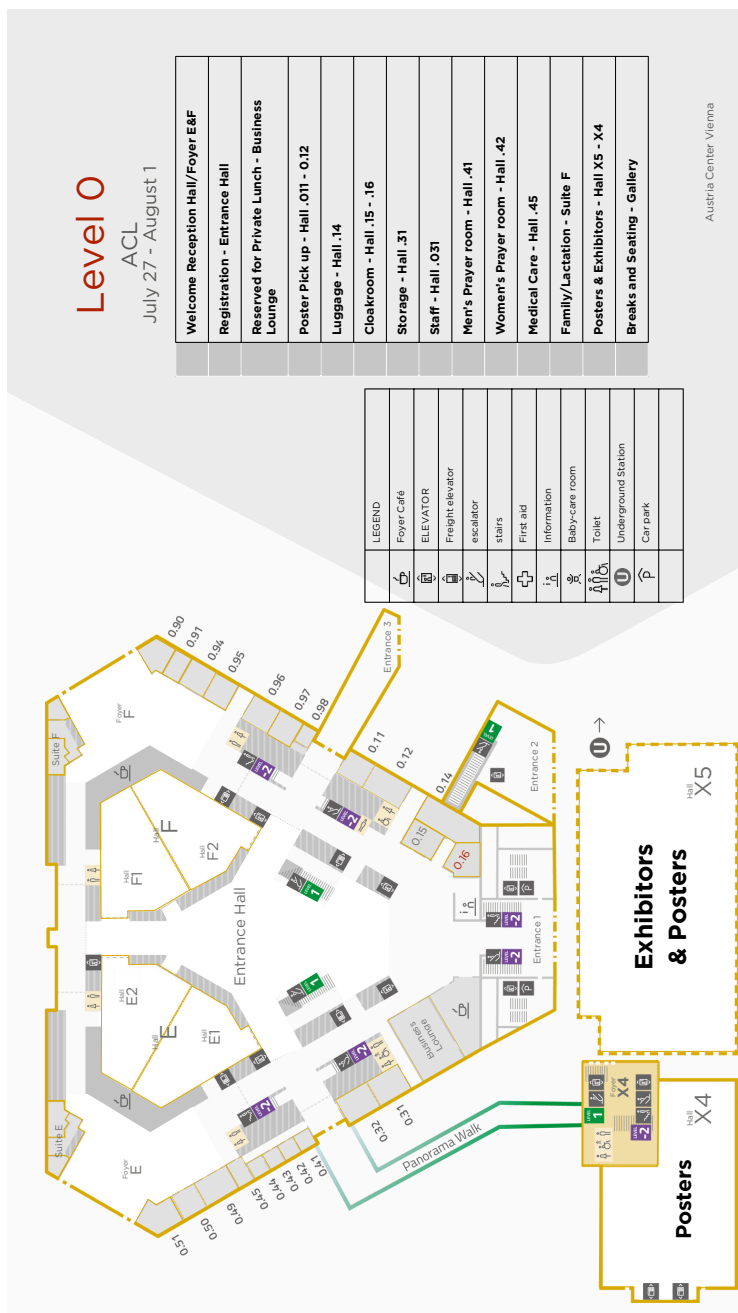
Popular Shopping Areas

- Mariahilfer Straße: <https://www.wien.info/de/sehen-erleben/shopping/shopping-rund-um-die-altstadt>
- Kärntner Straße: <https://www.wien.info/en/see-do/shopping/old-city>

12

Venue Map

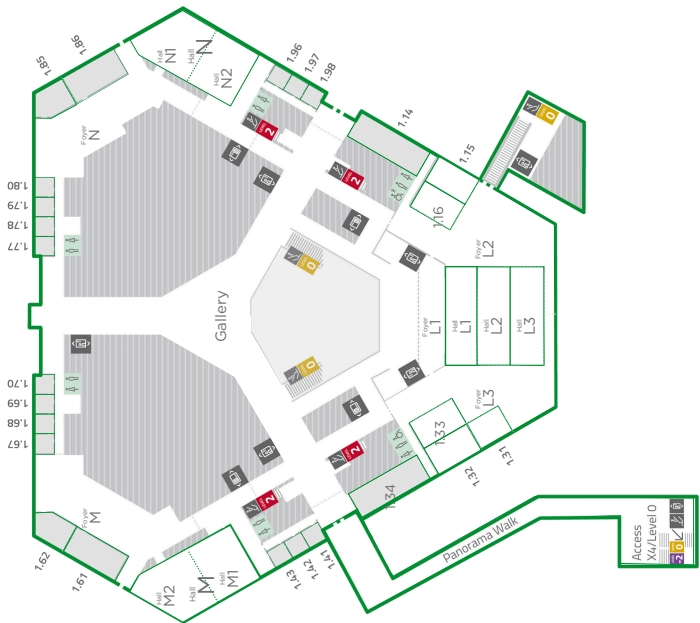




Level 1

ACL
July 27 - August 1

Tutorials (Sunday) Halls: M1-2, N1-2	
Orals (Mon-Wed) Halls: M1-2, N1-2, L2-3, 115-16, 161, 162, 185, 186	
Workshops (Thur & Fri) Halls: M1-2, N1-2, L2-3, 115-16, 161, 162, 185, 186	
Workshop Posters (Thur & Fri) Hall L1- L3	
Birds of a Feather Halls: 114, 131-32, 133	











Austria Center Vienna

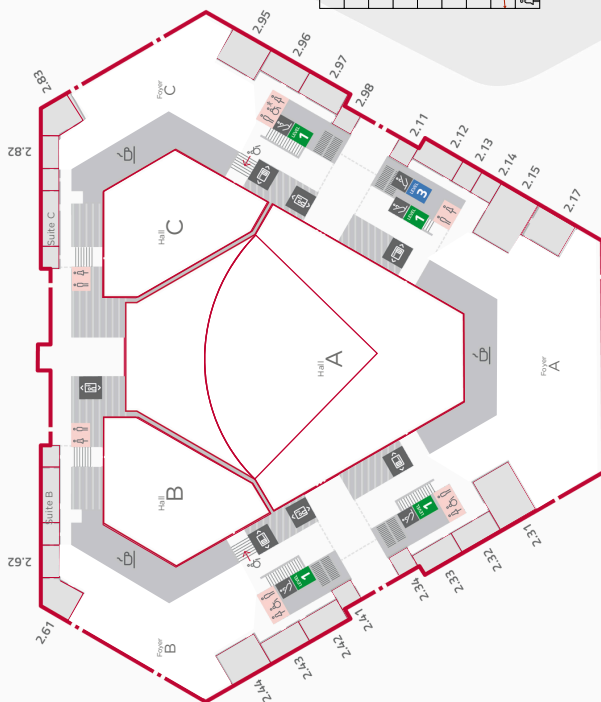
Level 2

ACL

July 27 - August 1

Tutorials (Sunday) Halls B & C
Plenary (Mon-Wed) Halls A-C
Panels/Orals (Wed) Halls A-C
Workshops (Thur & Fri) Halls: B 8C, 2.15, 2.17, 2.31, 2.44
Adhoc - Hall 2.12

	FOYER CAFÉ
	ELEVATOR
	Freight elevator
	escalator
	stairs
	Baby-care room
	Access for Wheelchairs
	Toilet



Austria Center Vienna

Alibaba Cloud Campus Recruitment

About Alibaba Cloud

Apsara Lab - Data to Intelligence Lab

Alibaba Cloud's Data to Intelligence Lab, led by world-renowned AI expert Prof. YE Jieping, is driving intelligent transformation across society with Qwen models. The team focuses on industry-specific LLMs, multimodal models, and foundational research on large models. Their innovations have been applied in various sectors, including smart traffic systems, urban governance, and biomedicine, while also enhancing major global sports events like the Asian Games and the Olympics.

Alibaba Cloud Computing Platform

Alibaba Cloud Computing Platform includes a proprietary big data platform ODPS (MaxCompute/Hologres/DataWorks, etc.), open-source big data platforms (E-MapReduce/Realtime Compute for Apache Flink, etc.), and an artificial intelligence platform PAI, etc. It provides a comprehensive product system covering data collection, storage and analysis, development and governance, AI training and application. Meanwhile, Alibaba Cloud Computing Platform is also equipped with a variety of computing capabilities including large-scale batch processing, real-time streaming processing, and machine learning. It can concurrently schedule over 100,000 servers for parallel computing. Alibaba Cloud is the only Chinese Company in Forrester's Leaders quadrant for cloud data warehouses. In terms of AI infrastructure, related products of Alibaba Cloud Computing Platform ranked second globally and first in Asia for product and technical capabilities in the Forrester Wave: AI Infrastructure Solutions 2024 evaluation.

</SHAPE
YOUR OWN FUTURE.

Application method



Scan the QR
code to submit
your resume

Target students

Currently enrolled
university students

Positions

Apsara Lab-Data to Intelligence Lab

- LLM Algorithm Research Intern

Alibaba Cloud Computing Platform

- Research on the application of large model downstream task transfer technology in general content security and agent security
- Research on knowledge distillation techniques for LLMs
- Research on key technologies for the pre-training of ultra-large-scale MoE models
- Exploration of agent-based adaptive RAG
- Optimization of large model hallucination in RAG scenarios

Locations: Hangzhou, Beijing

</SHAPE
YOUR OWN FUTURE.

ANT STAR⁺

Plan A⁺ Talent Program

Join Ant Group and become a member of "Plan A"
Light up the future with technology together!

"Plan A" is a new elite AI talent initiative under "Ant Star." It aims to recruit world-class artificial intelligence researchers from leading global universities to join Ant Group's exploration of AGI.

The letter "A" embodies two core ideals: Apex, Visionary pioneers reaching the summit, daring to push the boundaries of intelligence. Action, Pragmatic innovators committed to ensure AI benefits all humanity.

We discover and empower exceptional rising stars worldwide—passionate tech innovators with outstanding research enthusiasm. Through top-tier challenges, resources, and development spaces, we prepare tomorrow's AI pioneers.



Scan the code to
learn about Ant Group

The Minds We Seek

Top talents in LLM, Multimodal, Infra, RL, Diffusion, Sparse Attention, and Synthetic data.

Recruitment Position:

☒ Research Scientist -

Reinforcement Learning

Deep Learning

Synthetic Data

Diffusion

Inference

LLM Pre-training

LLM Post-training

Multimodal

What We Offer



Frontier Technology Arena

- Industry pioneer mentorship
- Mission-critical R&D challenges
- Strategic core tech domains



Unconstrained Innovation

- Personalized development roadmaps
- Full research freedom
- Unlimited GPU resources



Elite Growth Ecosystem

- Levels and salaries based on ability, not seniority
- Get recognized faster for what you achieve



Global R&D network

- Hubs across Hangzhou/Beijing/Shanghai/Singapore
- Integrated research labs
- Cross-disciplinary exchanges

Recruitment Target

Eligibility: Open to master's/PhD graduates in computer science, software engineering, AI, cybersecurity, ICT, mathematics, statistics, and interdisciplinary STEM fields.

Graduation Window: Candidates graduating in or after November 2025.

Join Us

Scan the QR code to complete the resume submission.

Contact Us:
AntTalent@antgroup.com



AI at Bloomberg

Bloomberg is building the world's most trusted information network for financial professionals. Our 9,000+ engineers, developers, and data scientists are dedicated to advancing and building new solutions and systems for the Bloomberg Terminal and other products in order to solve complex, real-world problems.

We have 15+ years of experience building and using AI solutions to help process and organize the ever-increasing volume of structured and unstructured financial information. In fact, we were one of the first on the street to use AI. Our team of 350+ AI researchers and engineers deploy different AI technologies, including ML, NLP, computer vision, information retrieval, and generative models, as well as neural networks and time-series analysis. This is how we provide our customers around the globe with the reliable, high-quality data, news, and information they need to make well-informed decisions and be more productive.

Learn more here:



#MAKEITHAPPENHERE

Bloomberg

Engineering



Make it
happen here.



CITADEL

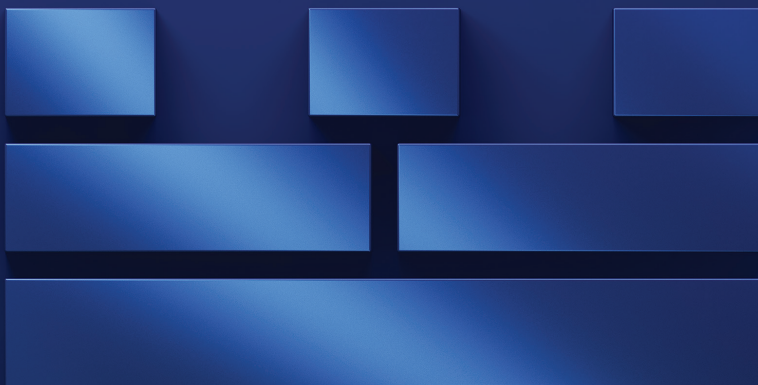


CITADEL | Securities

Build, test, and refine your ideas at the speed of the markets.

The markets present rich data, complex problems and fast feedback. We treat market problems as scientific questions. Our researchers can quickly see how the markets validate or refute their thinking.

Our culture of achievement invites the brightest minds in science. As a result, we have some of the most talented and accomplished researchers in the world working alongside our investment professionals, traders, and engineers.



Learn more at
events.citadel.com/acl2025

citadel.com | citadelsecurities.com



The all-in-one platform for
private and secure AI.

✂ @cohere

🔗 cohere.com

in cohere

✂ Cohere Labs

Exploring the unknown, together.
Changing spaces where
breakthroughs happen.

✂ @cohere_labs

🔗 cohere.com/research

in cohere labs



Where your skills meet
the right opportunity



Linked in



About Xiaomi

Xiaomi Corporation was founded in April 2010 and listed on the Main Board of the Hong Kong Stock Exchange on July 9, 2018 (1810.HK). Xiaomi is a consumer electronics and smart manufacturing company with smartphones and smart hardware connected by an IoT platform at its core.

Embracing our vision of "Make friends with users and be the coolest company in the users' hearts", Xiaomi continuously pursues innovations, high-quality user experience and operational efficiency. The company relentlessly builds amazing products with honest prices to let everyone in the world enjoy a better life through innovative technology.

Xiaomi Top Campus Talent Recruitment

100+cutting-edge research topics, orchestrating the future of technological innovation!

Covering frontier domains: Large Models, Multimodality, Machine Learning, Embodied Intelligence, Semiconductor Chips, Autonomous Driving, etc.



Elite Mentorship
from Top Industry
Experts



Premium Resource
Allocation for Cutting-
Edge Research



Boundless Career
Growth
Opportunities



Engagement in
Pioneering
Research Projects



Competitive
Compensation &
Benefits

Recruitment Targets

Internship Program:

Current undergraduate, master's or PhD students

Campus recruitment:

- Graduates graduating between 2023 and 2026 (inclusive)
- Postdoctoral researchers completing their programs in 2026

Base: Beijing | Shanghai | Wuhan



Join Us



**Pushing the boundaries of AI
through research, infrastructure
and product innovation.**

ACL 2025



Amazon Science



Academics at Amazon



Internships at Amazon



Come build the future with us

At Amazon, we believe that scientific innovation is essential to being the most customer-centric company in the world. Our scientists' ability to have an impact at scale allows us to attract some of the brightest minds in artificial intelligence and related fields.

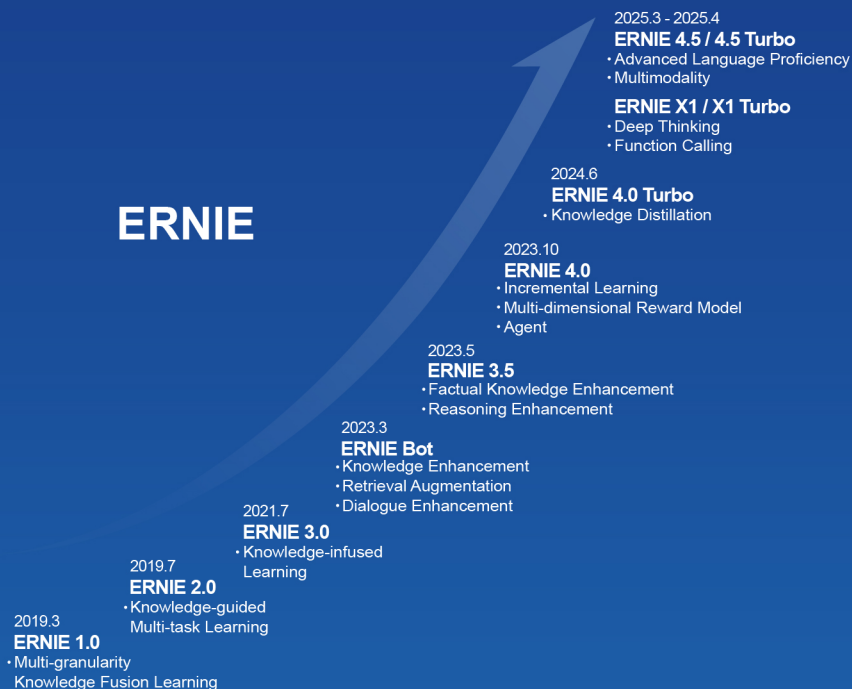




BAIDU WENXIN

A series of knowledge-enhanced large models, from general-purpose to industry-specific, independently developed by Baidu, speed up AI innovations and empower the industry to upgrade.

ERNIE



Knowledge Enhancement	Retrieval Augmentation	Dialogue Enhancement	Agents
Supervised Fine-Tuning	Reinforcement Learning	Prompt	
Pre-Training Models			

About Seed Team

The ByteDance Seed team was established in 2023, dedicated to discovering new approaches to general intelligence, pushing the boundaries of AI, and exploring new forms of interaction. The team's research areas include LLM, speech, vision, world models, infrastructure, AI Infra, and next-generation AI interactions. The team operates labs and positions in China, Singapore, the United States, and other locations.

With a long-term vision and determination in the AI field, the ByteDance Seed team remains committed to foundational research and aims to become a world-class AI research team, contributing to technological and societal advancement. The team has already launched industry-leading general large models and cutting-edge multimodal capabilities, supporting over 50 application scenarios such as Doubao, Coze, and Jimeng.

Jobs Information

Job Locations

Beijing, Shanghai, Hangzhou, Shenzhen, Singapore, San Jose, Seattle, and more.



Scan the QR code and check out the positions of Seed team



Follow Us

Visit the official website of the Seed team

 seed.bytedance.com

Stay updated on our latest research, model performance, and job opportunities by following us on **RED** and **Zhihu**









  Seed 研究员

  字节跳动 Seed

Scan the QR code to follow the official WeChat account



Founded in 2012, Noah's Ark Lab is the AI research centre for Huawei Technologies. We undertake the following primary domains: AI theory, AI data technology, AI system and engineering, efficient AI, computer vision, natural language processing, decision making & reasoning, recommendation & search, to drive fundamental R&D in AI.

 <i>AI Theory</i>	<i>Natural Language Processing</i> 
 <i>AI Data Technology</i>	<i>Computer Vision</i> 
 <i>AI System and Engineering</i>	<i>Decision Making & Reasoning</i> 
 <i>Efficient AI</i>	<i>Recommendation & Search</i> 

Mission

Our mission is to build practical AI technology to solve real world problems. We entail cutting-edge AI technologies including reinforcement learning, unsupervised learning, large foundation models, etc., in multiple areas, such as network intelligence (automatic control, security and maintenance), enterprise intelligence (e.g., safe city, smart supply chain, smart accounting) and terminal intelligence (e.g., AI camera, smartphones, and personalized recommendation).

Collaboration

We have top-tier researchers from world-leading universities and industry pioneers, collaborating with global innovative organizations including 40+ universities and 1,000+ researchers from 10+ countries. Through joint labs with 10+ top-tier AI universities, we pioneer breakthroughs in large language models, multimodal AI, agent systems, etc. Our achievements include multiple best/outstanding paper awards at top-tier conferences, building Huawei's competitiveness in AI.

Global Vision

Now, Noah's Ark Lab has established global footprint including Shenzhen, Beijing, Shanghai, Xi'an, Nanjing, Hefei, Hong Kong, Singapore, Paris, London, Montreal, Toronto, etc, and built global research partnership with over 12 top universities.

Join Us

We are hiring researchers in the related areas at all levels, including junior and senior positions. Interested candidates should send application materials including resume to noahlab1@huawei.com.





Inventing What's Next

IBM Research is home to 3,000 scientists and researchers who deeply believe in the power of the scientific method to invent what's next in computing for our company, for our clients and the world. We are pioneering technologies that will transform industries and society, including the future of AI, hybrid cloud, quantum computing and semiconductors. We've been here since the earliest days of computing, and we're leading the charge for the future. Since our first lab opened in 1945, we've authored more than 110,000 research publications. Our researchers have won six Nobel Prizes, six Turing Awards, and IBM has been granted more than 150,000 patents.

ibm.research.com



Megagon Labs

**Connecting AI, Data, and People
for Better Opportunities**

ACL Main:

CypherBench: Towards Precise Retrieval over Full-scale Modern Knowledge Graphs in the LLM Era

ACL Findings:

FactLens: Benchmarking Fine-Grained Fact Verification



blue

Research that ships.

Blue is an open-source agentic framework for building compound AI systems that integrate language models, data, APIs, and agents together for performing complex tasks. It's built for researchers exploring agentic systems, planning, data management, and data centric NLP in enterprise scale settings. Use Blue for your projects or contribute to it's development. Go to our Github →



ORACLE

Real careers in artificial intelligence

Oracle is proud to sponsor the 63rd Annual Meeting of the Association for Computational Linguistics. And we're just as proud to drive the future of AI, leveraging cutting-edge agentic platforms and multimodal GenAI to empower customers. Our teams are pushing the boundaries of innovation, from reasoning and planning to retrieval-augmented generation, GenAI for Databases and Analytics, NL2Code, and model architecture for large-scale training/inference.

Now, it's your turn. Push your career at the forefront of innovation and put your knowledge to work with the cloud leader for global business.

Take the next step in your career. Scan the QR code to explore opportunities with the Oracle Cloud Infrastructure AI Science team.



Create the future with us





WHERE DATA DOES MORE

Snowflake AI Data Cloud.
The platform for the AI era.

snowflake.com



TENCENT PROJECT UP

腾 讯 青 云 计 划

About

PROJECT UP

Project UP is a Tencent talent program with the aim of recruiting top technology students from around the world

We Connect - Campus Collaboration

A new collaboration model that enables you to join Tencent's core research projects on campus

We Discover - Global Premier Events

Expanded-scale events: UP Open Days, City Exclusives to Starlit Tech Gala

We Inspire - Expertise Meets Excellence

Engagement with Tencent's tech executives and distinguished scientists

Meet Us at
Booth 12



Tencent Project UP
Official Website



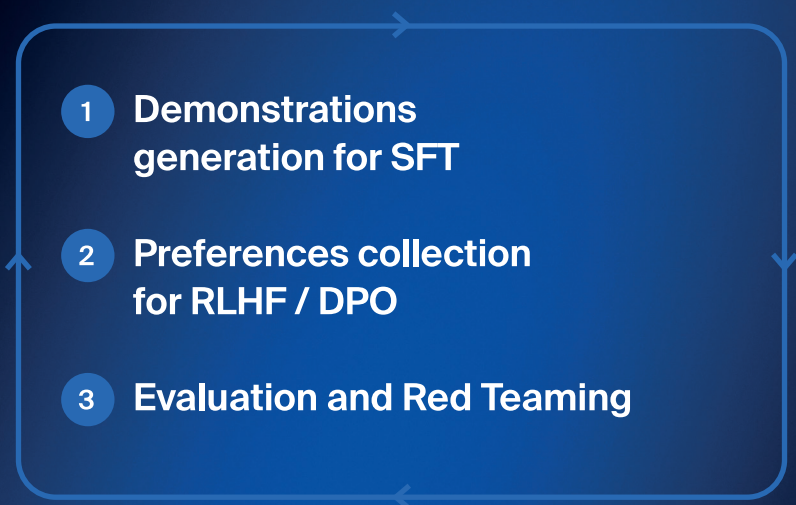
Tencent Project UP
Global Information



Connect with
our HR Team



Expertly crafted data for AI agents & models



All data types (text, image, video, audio)



Diverse languages



Specialized domains (coding, law, finance, more)



toloka.ai

LLM Research on the world's most exciting data: Machine Learning & NLP at Thomson Reuters

Train, evaluate, create data for, prototype with, and adversarially test models trained on data created by the world's most experienced journalists, editors, and lawyers. Research and experiment in collaboration with world-leading universities and publish at the best venues for LLM Research.



Join our team

© 2025 Thomson Reuters TR5806232/6-25



| About Kuaishou

Kuaishou is a leading global content community and social platform. Kuaishou has developed several cutting-edge models such as the Large Language Model "KwaiYii", Video Generation Large Model "Kling", Image Generation Large Model "Kolors", and Large Language Model for Coding "Kwaipilot". These models are deeply integrated into Kuaishou's core businesses, including short videos, live streaming, online advertising, e-commerce, and local services. They have also inspired innovative applications like "KlingAI", "AI Xiao Kuai", and "Kwaipilot".

| Join Us!

The KSTAR-X Talents Program is hiring now. Scan the QR code to learn more and submit your application, or reach out via campus@kuaishou.com





APPLY FOR OUR AI ACCELERATED RESEARCH PRIZE

Showcase your innovative work, gain global recognition, and connect with leading AI pioneers.

cheninstitute.org/prize

Appen

Build Safer, Smarter AI with High-Quality, Bespoke Data

For over 25 years, Appen has enabled the world's top AI teams to fine-tune and evaluate models using high-quality, human-labeled data across 500+ languages and 200+ countries.

From multilingual LLMs to complex multimodal pipelines, our experts deliver the precision your models need.

Visit us at Booth 8 appen.com



Adobe supports

ACL 2025

Adobe Research is Hiring!

Our Research Areas



Artificial Intelligence
& Machine Learning



Computer Vision,
Imaging & Video



Graphics (2D & 3D)



Content Intelligence



Audio



Document
Intelligence



AR, VR & 360
Photography



Natural Language
Processing



Human-Computer
Interaction



Data
Intelligence



Intelligent Agents
& Assistants



Systems &
Languages

If you're interested in research opportunities at Adobe scan the QR code to submit your contact information and resume.

Learn more:

Interns: [Research Adobe Express Page](#)

Experience: research.adobe.com



[Register here or scan the QR code](#)



Language generation, semantically grounded

Explore our technologies in detail
just scan the QR code below



<https://babelscape.com>

DEFINED.AI

Empower
your LLM



Discover powerful,
responsible, and
high-quality AI.
Smarter fine-tuning.
Safer data.
Superior models.

Find out more



JPMorganChase

MACHINE LEARNING
CENTER OF EXCELLENCE

Meet the MLCOE, leaders in machine learning
and innovation. Visit us to learn how our
diverse team is delivering AI/ML solutions to
transform finance!

We're looking for curious problem-solvers
with a passion for developing innovative ML
solutions.

See you at
Booth 5!

jpmorgan.com/mlcoe



NEC

\Orchestrating a brighter world

Advancing agentic
AI for greater social value.



#NECLabs
neclab.eu





\$100,000

to fund language technology innovators who share the goal of making it easier for everyone to understand and be understood by all others.

Find out more.



DataoceanAI

Empower Your AI with Best Data



OFF-THE-SHELF DATASETS



DATA SERVICES



DOTS PLATFORM



MODEL EVALUATION



CONTACT US

Dataocean AI Inc.
100 N Howard St Ste R, Spokane, WA 99201
contact@dataoceanai.com
www.dataoceanai.com