ACL 2026

The 64th Annual Meeting of the Association for Computational Linguistics (ACL 2026)

Tutorial Abstracts

July 2-7, 2026

The ACL organizers gratefully acknowledge the support from the following sponsors.

Diamond

Platinum


Gold

Silver

Bronze

Academic Ally

Order copies of this and other ACL proceedings from:

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

ISBN 979-8-89176-394-4

Introduction

We are excited to welcome you to the tutorial session of ACL 2026!

Our tutorials cover a range of topics that have moved to the forefront of NLP research in 2026. They cover diverse areas in core engineering research such as data, agents, and reasoning, as well as dealing with broader fields such as the future of work and bridging academia, industry, and society. While this is only a small subset of topics that are of potential interest to the broader ACL community, we hope this provides an interesting sample of themes you will see in the broader conference and research landscape.

This year, we did a joint call for proposals with EACL alongside their tutorial co-chairs: Chenghua Lin, Aline Paes, and Rodrigo Wilkens. We received 61 submissions of which 4 were accepted to EACL and 6 are appearing here at ACL. All of the proposals were reviewed by multiple of the tutorial co-chairs before a broader discussion among all. As with previous years, it was a highly competitive selection process with many excellent submissions. While 6 tutorials by no means capture all of the potentially interesting topics relevant to our research community, we anticipate that these will engage attendees from a wide range of backgrounds and expertise levels to kick-start the first day of ACL 2026.

We very much appreciate the efforts of all the tutorial authors who provide compelling lectures and talks and make this day possible. We would also like to extend a special thanks to Jenn Rachford who works tirelessly behind the scenes to help organize all parts of the conference - and the tutorials are no exception.

We hope you enjoy the conference,
Jacob Andreas and Kenton Murray
ACL 2026 Tutorial Co-Chairs

Organizing Committee

General Chair

Philipp Koehn, Johns Hopkins University

Program Chairs

Maria Liakata, Queen Mary University of London
Viviane P. Moreira, Institute of Informatics at UFRGS, Brasil
Jiajun Zhang, Chinese Academy of Sciences
David Jurgens, University of Michigan

Tutorial Chairs

Jacob Andreas, Massachusetts Institute of Technology
Kenton Murray, Johns Hopkins University

Table of Contents

Program

Thursday, July 2, 2026

09:00 - 12:30 *Future of Work in the Age of LLMs*

09:00 - 12:30 *The Data Frontier for Large Language Models: Selection, Synthesis, and Tools*

09:00 - 12:30 *Towards Effective and Efficient Multi-Agent Language Model Systems: Foundations, Prospects, and Applications*

14:00 - 17:30 *Current Advances in LLM Reasoning*

14:00 - 17:30 *Knowledge Control for Responsible Generative AI: Bridging Academia, Industry, and Society*

14:00 - 17:30 *The Interplay between Metaphors and NLP*

Future of Work in the Age of LLMs

Zora Wang, Yijia Shao, David Nguyen, Diyi Yang

https://future-of-work-llm-tutorial.github.io/

The recent development of large language models (LLMs) has revolutionized the landscape of human work. These models possess the ability to follow complex human instructions and operate versatile computer software, enabling them to participate in, augment, or even automate realistic occupational tasks that once thought to be exclusive to humans. As LLMs are increasingly integrated into workplaces, they are already reshaping labor dynamics (Hoffmann et al., 2024; Demirci et al., 2025) and raising urgent concerns about job displacement, diminished human agency, and overreliance on automation (Hazra et al., 2025). As a result, the future of work is undergoing a profound transformation: How will human occupations and task requirements evolve over time? And what roles will LLM-based systems play, as they become increasingly capable collaborators and autonomous workers? And how can we build technological and data infrastructures to support human-AI collaboration? This tutorial will provide an overview of the future of work shaped by the interplay of LLMs and humans, examining the emerging challenges, opportunities, and ethical considerations in this dynamic landscape. We begin by examining the economic landscape of work and how NLP technologies drive automation, followed by methods for developing LLMs that augment human labor and recent advances in LLM-based agents. We then cover evaluation approaches for workplace contexts, including datasets, benchmarks, and metrics, and conclude with open questions on technical, human, and societal implications.

Zora Wang
email: zhiruow@cs.cmu.edu
website: https://zorazrw.github.io/
Zora Zhiruo Wang is a PhD student at Carnegie Mellon University, Language Technologies Institute, advised by Professors Graham Neubig and Daniel Fried. Her research focuses on using programmatic approaches to solve real-world problems, especially those related to human occupations. Zora has organized Deep Learning for Code workshops at ICLR, NeurIPS, and ICML, given a tutorial about LLMs for tabular data at SIGIR 2024, and presented her work at and served as reviewers for top-tier NLP/ML conferences including NeurIPS, ICLR, ACL, EMNLP, and KDD. She has been recognized with Google PhD Fellowship and CMU Presidential Fellowship.

Yijia Shao
email: shaoyj@stanford.edu
website: https://cs.stanford.edu/~shaoyj/
Yijia Shao is a PhD student at Stanford University, advised by Professor Diyi Yang. Her research focuses on human-agent collaboration, specifically positioning NLP models into broader systems (e.g. LLM agents) and optimizing them to collaborate with humans. Her work appears in top ML and NLP conferences, including NeurIPS, ICLR, ACL, EMNLP, and NAACL. She is supported by Stanford School of Engineering Fellowship and Thinking Machines Fellowship.

David Nguyen
email: davidngu@stanford.edu
website: https://digitaleconomy.stanford.edu/person/david-nguyen/
David Nguyen is a Research Scientist at the Stanford Digital Economy Lab within the Institute for

Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 5: Tutorial Abstracts), pages 1–2

Human-Centered Artificial Intelligence (HAI), and a Research Associate at the Economic Statistics Centre of Excellence (ESCoE). His research explores new ways to analyze and measure changing, modern, and digital economies. He is particularly interested in advancing economic indicators and statistics on economic output and welfare.

Diyi Yang

email: diyiy@cs.stanford.edu

website: https://cs.stanford.edu/~diyiy/

Diyi Yang is an assistant professor in the Computer Science Department at Stanford University. Her research focuses on human-centered natural language processing and computational social science. Diyi has organized four workshops and seven tutorials at *CL conferences, with recent ones including "Language Agents: Foundations, Prospects, and Risks" at EMNLP 2024, "Social Intelligence in the Age of LLMs" at NAACL 2025, and "Human-AI Collaboration: How AIs Augment Human Teammates" at ACL 2025.

2

The Data Frontier for Large Language Models: Selection, Synthesis, and Tools

Lijun Wu, Wentao Zhang, Conghui He

As the development of Large Language Models (LLMs) matures, the focus of the research community is undergoing a critical shift from a purely model-centric to a data-centric paradigm. It is now evident that the quality, diversity, and composition of training data—not merely its scale—are the primary drivers of a model's advanced capabilities, from complex reasoning to reliable instruction following. However, acquiring and curating such high-quality data remains a significant bottleneck. This tutorial provides a comprehensive and practical guide to the state-of-the-art in data research directions for LLMs. We structure the tutorial around the two core pillars of modern data strategy: intelligent data selection and advanced data synthesis. In the first part, we delve into methods for curating the most valuable information from vast, noisy datasets, covering techniques like LLM-as-a-judge for automated quality filtering and active learning for maximizing annotation efficiency. The second part explores the synthetic data revolution, detailing paradigms that range from generating complex reasoning traces (e.g., Chain-of-Thought) to deploying sophisticated multi-agent workflows that can autonomously create high-quality, diverse instruction data from raw seeds. Finally, we will conclude with a practical overview of open-source tools and platforms that facilitate these data-centric workflows, empowering researchers and practitioners to build better models through better data. Attendees will leave with a principled framework and actionable insights for designing and implementing the advanced data strategies required to build the next generation of powerful, specialized, and aligned LLMs.

Lijun Wu
email: lijun_wu@outlook.com
website: https://apeterswu.github.io/
Dr. Lijun Wu is a Young Scientist in Shanghai AI Laboratory. Previously, he was a Research Scientist in ByteDance, a Senior Researcher in Microsoft Research. He got the Ph.D. degree from Sun Yat-sen University (SYSU), and was a member of joint Ph.D. program between SYSU and MSRA. His research interests are on AI/LLMs (e.g., data-centric intelligence, SFT/RL), AI4Science (e.g., LLM4Science, scientific reasoning).

Wentao Zhang
email: weantao.zhang@pku.edu.cn
website: https://zwt233.github.io/
Dr. Wentao Zhang is an assistant professor (Principal Investigator/PhD Advisor) in the Center of Machine Learning Research at Peking University (PKU), and he leads the Data-centric Artificial Intelligence (DCAI) group. Wentao's research focuses on DCAI, LLM, AI systems and AI4Science. Wentao is the contributor or designer of several system projects, including DataFlow, MinerU, and Angel. Before joining PKU, wen- tao worked as a research fellow with Prof. Jian Tang at Montreal Institute for Learning Algo- rithms (Mila, led by Prof.Yoshua Bengio), and he received his Ph.D. degree in CS at PKU, supervised by Prof. Bin Cui.

Conghui He
email: heconghui@pjlab.org.cn
website: https://conghui.ai/

Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 5: Tutorial Abstracts), pages 3–4

Dr. Conghui He is a Young Leading Scientist at the Shanghai AI Lab and an Adjunct Doctoral Supervisor at School Shanghai Jiao Tong University. Recognized as a National-level Young Talent, he hold a Ph.D. from Tsinghua University and was a visiting researcher at Stanford University and Imperial College London. He is the creator of MinerU, the world's leading open-source data engine for large models. Additionally, he oversees a dedicated data team that curates high-quality datasets for leading models such as InternLM and InternVL.

Towards Effective and Efficient Multi-Agent Language Model Systems: Foundations, Prospects, and Applications

Xuan Wang, Shuxiang Cao, Yuchen Zhuang, Wenqi Shi

https://xuanwang91.github.io/2026-07-02-acl26-tutorial

Multi-agent systems powered by large language models (LLMs) offer a promising paradigm for tackling complex reasoning, decision-making, and problem-solving tasks. However, achieving both effectiveness and efficiency in such systems remains a critical challenge. This tutorial introduces recent advances in building effective and efficient multi-agent LLM systems, focusing on three core components. First, we discuss the design of individual LLM agents. We present state-of-the-art techniques for enabling capable agents using efficient and compact LLMs, including model distillation, dynamic routing, and memory- and compute efficient serving, providing a foundation for scalable and responsive agent design under resource constraints. Second, we cover coordination and communication among agents, crucial for collective performance, highlighting methods for improving multi-agent reasoning and decision-making through prompt and graph optimization, sycophancy mitigation, and structured LLM-based frameworks. Last, we explore real-world applications of LLM agents in areas such as industry, healthcare, quantum computing, and various scientific domains.

Xuan Wang

email: xuanw@vt.edu

website: https://xuanwang91.github.io/

Dr. Xuan Wang is an Assistant Professor in the Department of Computer Science at Virginia Tech. Her research interests are in natural language processing, data mining, AI for sciences, and AI for healthcare. Xuan was a recipient of the NSF CAEER Award 2025, Cisco Research Award 2025, NSF NAIRR Pilot Award 2024-2025, and NAACL Best Demo Paper Award 2021. Xuan has served as a Program Chair of the SouthNLP Symposium 2024 (>150 participants), from more than 20 universities across the USA. She has also served as a Program Chair for the Undergraduate and High School Symposium at IEEEBigData 2024 and IEEE-ICDM 2025. She has also served as a Senior Area Chair, Area Chair, and Program Committee in major AI conferences (e.g., ARR, ACL, EMNLP, NAACL, NeurIPS, ICLR, KDD). Xuan has delivered tutorials in AAAI 2025, EMNLP 2024, KDD 2022, TheWebConf 2022, and IEEE-BigData 2019.

Shuxiang Cao

email: shuxiangc@nvidia.com

website: https://www.physics.ox.ac.uk/our-people/caos

Dr. Shuxiang Cao is a Senior Research Scientist at NVIDIA and a Long Term Visitor in the Department of Physics at the University of Oxford. His research is at the intersection of quantum technology, artificial intelligence, and high-performance computing. He obtained his PhD in Physics (Quantum Technology) from the University of Oxford, where he worked on superconducting quantum computers. At NVIDIA, his work focuses on advancing artificial intelligence for science, including the development of large language model applications for scientific research and autonomous laboratories that accelerate scientific discovery. Dr. Cao has published in high profile journals in physics, including Physical Review Letters, Physics Reports, and Science Advances, as well as at computer science conferences, including NeurIPS, ACL, and EMNLP.

Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 5: Tutorial Abstracts), pages 5–6

Yuchen Zhuang

email: yczhuang@google.com

website: https://night-chen.github.io/

Dr. Yuchen Zhuang is a Research Scientist at Google DeepMind. His research focuses on language intelligence, with the goal of developing LLM-based agents that can show human-like reasoning and planning on challenging real-world tasks. His experience covers model pre-training, instruction fine-tuning, reinforcement learning from human feedback (RLHF), and evaluation. He has published more than 30 papers in top AI venues such as NeurIPS, ICLR, ICML, ACL, EMNLP, NAACL, UAI, and KDD. He has received the 2024 J.P. Morgan PhD Fellowship Award, the NeurIPS 2023 Scholar Award, and the Best SIGBio Paper Award at ACM BCB. He serves as an Area Chair and Program Committee member in major AI conferences (e.g., ARR, ACL, EMNLP, NAACL, NeurIPS, ICLR, ICML).

Wenqi Shi

email: wenqi.shi@utsouthwestern.edu

website: https://wshi83.github.io/

Dr. Wenqi Shi is an Assistant Professor in the Department of Health Data Science and Biostatistics at the University of Texas Southwestern Medical Center. Her research focuses on the intersection of AI and healthcare, advancing both foundational algorithms and agentic systems for precision and personalized medicine. Her recent work includes developing LLMs for translational medicine, advancing agentic AI for biomedical discovery, and promoting responsible AI practices to improve clinical research and care. She has received the NVIDIA Academic Grant Program Award (2025), the Texas Advanced Computing Center (TACC) Pilot Award (2025), the Best SIGBio Paper Award at ACM BCB (2023), and recognition as a Rising Star in EECS (2023). She has served as Finance Co-Chair of ACM BCB 2025, Publication Co-Chair of IEEE BHI 2025 (both >300 participants), and as an Area Chair or Program Committee member in major AI and healthcare conferences (e.g., ACL, EMNLP, NAACL, NeurIPS, ICLR, ICML, BCB, BIBM).

Current Advances in LLM Reasoning

Akhil Arora, Vishrav Chaudhary, Julia Kreutzer, Nearchos Potamitis, Nouha Dziri, Niket Tandon

As large language models (LLMs) increasingly tackle reasoning-heavy tasks, from mathematics to commonsense to multilingual understanding, researchers face three pressing questions: How well do models reason? How can we make them reason better? What are the next frontiers in LLM reasoning? This tutorial answers these questions through a unified view of LLM reasoning. This tutorial explores comprehensive evaluation strategies to assess the reasoning abilities of models and discusses two types of methods to improve models' reasoning: advanced inference time methods, such as structured and self-improvement inference methods, and (ii) post-training methods, such as RLHF, DPO, and GRPO that aim to make LLMs think more like humans. The tutorial explores these technical discussions while maintaining a practical outlook through illustrative demos and short guided hands-on exercises. The tutorial is designed for both researchers and practitioners seeking practical insights into LLM reasoning.

Akhil Arora
email: akhil.arora@cs.au.dk
website: https://au-clan.github.io/people/aarora/
Akhil Arora is a professor of computer science at Aarhus University (AU), where he leads the CLAN for AI Research on Language and Networks (or "CLAN" for short). He is a fellow of the Copenhagen Center for Social Data Science, an affiliate of the Pioneer Centre for AI, and a formal collaborator of the Wikimedia Foundation. Akhil's research lies broadly in human-centered AI with a focus on improving human knowledge-seeking, bridging knowledge gaps, and promoting knowledge equity on the Web.

Vishrav Chaudhary
email: vishrav@meta.com
Vishrav Chaudhary is a Research Scientist working on post-training research and reasoning capabilities in large and small LMs. His research spans efficient multilingual modeling, reinforcement learning from human feedback, and cross-lingual reasoning. Vishrav has over a decade of experience leading major multilingual modeling efforts at Meta, Microsoft Turing, and FAIR, including contributions to the LLaMA-4, Phi3, and Phi4 model families.

Julia Kreutzer
email: juliakreutzer@cohere.com
website: https://juliakreutzer.github.io/
Julia Kreutzer is a Senior Research Scientist focusing on multilingual reasoning and evaluation for LLMs. She has previously worked at Google Translate and obtained her PhD from the University of Heidelberg. Julia has led research at the intersection of multilingual evaluation, post-training, and inclusivity in language technologies.

Nearchos Potamitis
email: nearchos.potamitis@cs.au.dk
website: https://potamitisn.github.io/
Nearchos Potamitis is a PhD student in Computer Science at Aarhus University (AU), advised by Akhil Arora. His research explores reasoning, efficiency, and evaluation in LMs, focusing on practical

Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 5: Tutorial Abstracts), pages 7–8

systems for robust and affordable inference. Nearchos has prior experience developing educational materials and interactive demos for graduate NLP courses at AU and for inter- national summer schools.

Nouha Dziri

email: nouha.dziri@gmail.com

website: https://nouhadziri.github.io/

Nouha Dziri is a Research Scientist at the Allen Institute for AI (AI2) working with Yejin Choi. Her research spans post-training, reasoning robustness, and LLM interpretability. Nouha has co-led the OLMo safety and post-training teams and previously worked at Google Research, Microsoft Research, and Mila.

Niket Tandon

email: nikett@gmail.com

website: https://niket.tandon.info/

Niket Tandon is a Principal Research Scientist focusing on feedback- guided and self-reflective reasoning in LLMs. Niket previously led reasoning research at AI2's Aristo team and earned his PhD from MPI for Informatics under Gerhard Weikum. His work has appeared in ACL, NeurIPS, and EMNLP, and includes the creation of large-scale commonsense resources like WebChild.

Knowledge Control for Responsible Generative AI: Bridging Academia, Industry, and Society

Zheyuan Liu, Yixin Wan, Kai-Wei Chang, Meng Jiang, Jieyu Zhao, Nouha Dziri, Yuning Mao, Jia-Chen Gu, Jindong Gu

https://franciscoliu.github.io/knowledge-control-tutorial-2026/

Controlling the knowledge and behavior of generative AI systems, including large language models (LLMs), multimodal LLMs (MLLMs), and text-to-image (T2I) models, has become critical as they are increasingly used in safety-sensitive and socially impactful applications. These models often encode unintended, biased, or private content, leading to harmful or unethical outputs. Post-training knowledge control has thus emerged as a practical framework for selectively modifying or removing model behaviors without full retraining, offering scalable and interpretable interventions for improving safety, privacy, and fairness. This tutorial introduces the foundations of post-training knowledge control and showcases recent frontier methods, bridging research insights with real-world practices from both academia and industry. We cover: (i) key motivations and failure modes, such as harmful generation and stereotype reinforcement; (ii) core methods such as machine unlearning, knowledge editing, and inference-time interventions for targeted behavior adjustment; and (iii) evaluation protocols for balancing forgetting, retention, and fairness. Case studies will span text and vision–language generation, including privacy preservation, bias mitigation, and factual correction.

Zheyuan Liu
email: zliu29@nd.edu
website: https://franciscoliu.github.io/
Zheyuan Liu is a Ph.D. candidate in Computer Science and Engineering at the University of Notre Dame, advised by Prof. Meng Jiang. His research focuses on trustworthy generative AI, including LLM/MLLM safety, AI privacy, fairness, agentic safety, and multimodal models. He received his B.S. degrees in Computer Science and Applied Mathematics from Brandeis University.

Yixin Wan
email: elaine1wan@g.ucla.edu
website: https://scholar.google.com/citations?user=hZPIICQAAAAJ&hl=en
Yixin Wan is a Ph.D. Candidate in Computer Science at the University of California, Los Angeles, advised by Prof. Kai-Wei Chang. Her research focuses on trustworthy NLP and multimodal generative models, including fairness, safety, text-to-image generation, vision-language models, and machine unlearning. She received her undergraduate training in mathematics at UCLA.

Kai-Wei Chang
email: kwchang@cs.ucla.edu
website: https://web.cs.ucla.edu/~kwchang/
Kai-Wei Chang is a Professor of Computer Science at the University of California, Los Angeles. He co-directs the UCLA DataX AI Technology Center and leads research on trustworthy AI, multimodal foundation models, reasoning, controllable generation, and human-centered AI systems. His work spans fairness, robustness, unlearning, safety, and reliable NLP and multimodal systems.

Meng Jiang
email: mjiang2@nd.edu

Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 5: Tutorial Abstracts), pages 9–10

website: http://www.meng-jiang.com/
Meng Jiang is the Frank M. Freimann Collegiate Professor of Computer Science and Engineering at the University of Notre Dame. His research interests include data mining, machine learning, natural language processing, artificial intelligence, and trustworthy generative AI. He has worked on topics such as computational behavior modeling, personalized language models, and machine unlearning.

Jieyu Zhao
email: jieyuz@usc.edu
website: https://jieyuz.net/
Jieyu Zhao is a Gabilan Assistant Professor in the Thomas Lord Department of Computer Science at the University of Southern California. She leads the LIME Lab, and her research focuses on trustworthy language models and human-centered AI. Before joining USC, she was an NSF Computing Innovation Fellow at the University of Maryland, College Park, and she received her Ph.D. in Computer Science from UCLA.

Nouha Dziri
email: nouha.dziri@gmail.com
website: https://nouhadziri.github.io/
Nouha Dziri is a Senior Research Scientist at Cohere Labs, after previously working as a research scientist and postdoctoral researcher at the Allen Institute for AI. Her research focuses on NLP and machine learning, with an emphasis on large language models, reasoning, safety, security, and post-training methods. Her work includes studies of hallucination, compositional reasoning, red-teaming, and safety-oriented post-training for open language models.

Yuning Mao
email: yuningm@meta.com
website: https://morningmoni.github.io/
Yuning Mao is a Research Scientist at Meta Superintelligence Labs. His work focuses on large language models, post-training, code generation, and helping users acquire information and knowledge more effectively. He received his Ph.D. in Computer Science from the University of Illinois Urbana-Champaign and has contributed to Meta's Llama series and related post-training efforts.

Jia-Chen Gu
email: gujc@ucla.edu
website: https://jasonforjoy.github.io/
Jia-Chen Gu is a Postdoctoral Researcher in the Department of Computer Science at the University of California, Los Angeles. He is hosted by Prof. Nanyun Peng and Prof. Kai-Wei Chang, and he received his Ph.D. from the University of Science and Technology of China in 2022. His research focuses on machine learning for natural language processing, especially retrieval-augmented language models, model editing, representation learning, and robust, efficient, and trustworthy AI.

Jindong Gu
email: jindong.gu@outlook.com
website: https://jindonggu.github.io/
Jindong Gu is a Senior Research Scientist at Google and a Senior Research Fellow at the University of Oxford. His research focuses on responsible AI, including the reliability, robustness, privacy, interpretability, and safety of AI systems. He received his Ph.D. from the University of Munich in 2022 and has worked on the safety of foundation models and the robustness of AI agents.

The Interplay between Metaphors and NLP

Joanne Boisson, Jose Camacho-Collados, Elisa Sanchez-Bayona, Rodrigo Agerri

https://sites.google.com/view/metaphor-nlp-tutorial/home

This tutorial will be of the type Introductory to CL/NLP topics, providing an overview of the metaphor processing field. The tutorial will discuss the influence of various metaphor theories on the creation of annotated resources and models. We will particularly focus on recent directions opened by LLMs for metaphor interpretation in multilingual and multimodal settings. Two types of audience may benefit from this tutorial: researchers in the humanities and computational social sciences interested in automatic or semi automatic metaphor analysis, and NLP researchers interested in understanding metaphor processing and improving metaphor modeling.

Joanne Boisson
email: boissonjc@cardiff.ac.uk
website: www.metaphorshare.com
Joanne Boisson is a postdoctoral researcher at Cardiff University working with Prof. Jose Camacho-Collados. She is completing her PhD on metaphor processing within the Cardiff Natural Language Processing group (Cardiff NLP) at the School of Computer Science, where her work focuses on developing resources for computational metaphor studies and on the identification of metaphors and proportional analogies. Prior to her doctoral research, she spent four years as a machine learning engineer and advisor specialising in Mandarin Chinese information extraction.

Jose Camacho-Collados
email: camachocolladosj@cardiff.ac.uk
website: https://josecamachocollados.com/
Jose Camacho-Collados is a UKRI Future Leaders Fellow and Professor at the School of Computer Science of Cardiff University, where he co-founded the Cardiff Natural Language Processing group (Cardiff NLP). Before joining Cardiff University, he completed his PhD in Sapienza University of Rome and was a Google AI PhD Fellow. Jose has worked in multiple NLP areas with a particular focus on semantics, multilinguality and computational social science with an interdisciplinary perspective.

Elisa Sanchez-Bayona
email: elisa.sanchez@ehu.eus
Elisa Sanchez-Bayona is a PhD student at the University of the Basque Country (EHU), currently developing her thesis on automatic metaphor processing from a cross-lingual perspective. Her contributions to the field include the creation of resources for Spanish metaphor detection and a parallel dataset for metaphor detection and interpretation tasks. Her work has been published in venues like ConLL, ACL and the Computational Linguistics journal.

Rodrigo Agerri
email: rodrigo.agerri@ehu.eus
website: https://ragerri.github.io/
Rodrigo Agerri is a permanent researcher at the HiTZ Center, University of the Basque Country EHU, and a member of ELLIS. Apart from investing a lot of time in metaphor-related research, he is generally interested in multilingual and crosslingual NLP with LLMs. He publishes in major venues, including

Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 5: Tutorial Abstracts), pages 11–12

Artificial Intelligence, Computational Linguistics, Artificial Intelligence in Medicine, ACL, or EMNLP. He has helped to organize several workshops/shared tasks such as BLEND SemEval 2026, Multilingual CounterSpeech Generation at COLING 2025, and ArgNLE at ECAI 2024.

Author Index