

ACL 2026

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

**Proceedings of the Conference - Volume 4: Student Research  
Workshop**

July 2-7, 2026

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

## Diamond

**Bloomberg**  
Engineering



## Platinum



## Gold



## Silver

**Adobe**

## Bronze

 translated.



The logo for ARC, consisting of the letters "ARC" in a black serif font with a red arc above the "A".

The logo for Abelscape, featuring a blue square icon with a white shape inside, followed by the text "abelscape®" in a blue sans-serif font.

The logo for DataMundi.ai, with "DATA" in blue, "mundi" in grey, and ".ai" in blue.

**NEC**

## Academic Ally



**W**  
MASTER OF SCIENCE IN  
COMPUTATIONAL LINGUISTICS  
UNIVERSITY of WASHINGTON

©2026 Association for Computational Linguistics

Order copies of this and other ACL proceedings from:

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

ISBN 979-8-89176-393-7

## Introduction

Welcome to the ACL 2026 Student Research Workshop (SRW)!

The Student Research Workshop continues its tradition of serving as a dedicated venue for student researchers in computational linguistics, natural language processing, and machine learning. Held in conjunction with ACL 2026 in San Diego, California, the SRW offers students an opportunity to present their work, engage with the broader research community, and receive constructive feedback from experienced researchers.

As in previous editions, the workshop welcomes two types of submissions: research papers and thesis proposals. Research papers may describe completed work or ongoing projects with preliminary results, while thesis proposals provide students with a platform to seek feedback on their broader research plans. The SRW also supports both archival and non-archival submissions, allowing students to share work at different stages of maturity while preserving flexibility for future publication. All submissions, regardless of track, are reviewed with the same care and commitment to mentorship.

This year, ACL 2026 SRW received a total of 402 submissions (379 direct submissions and 23 ARR commitments). Of these, we have accepted a total of 130 papers. Accepted papers will be presented during the conference with twelve featured as oral presentations. The final acceptance rate stands at 34.7%, comparable to last year's rate (32.2%). Of the accepted papers, 109 are archival and 21 are non-archival. Four papers were withdrawn after acceptance.

Mentorship remains central to the mission of the SRW. This year, 114 submissions participated in the pre-submission mentoring program, through which student authors received early feedback on their writing, presentation, and research framing from experienced mentors. We sincerely thank all 114 mentors for generously contributing their time and expertise throughout the submission, review, and revision process.

The student research workshop will be held on July 5th and 6th, for oral and poster presentations. For the virtual component of the conference, we keep as much as possible the spirit of an in-person conference. All talks and posters are pre-recorded and made available at the beginning of the conference for participants to watch asynchronously. Our oral session contains 12 talks followed by a Live QA part with the presenters. Topics span a broad range, with a strong emphasis on large language models, alongside interpretability, reasoning, evaluation, multilingual and low-resource NLP, uncertainty quantification, safety, and fairness, parameter-efficient methods and model efficiency, multi-agent systems, natural language inference, text classification, and clinical and legal applications.

This year's workshop has secured substantial funding to support student participation. Thanks to the dedicated efforts of the SRW faculty advisors, multiple sources of financial support were obtained. The faculty advisors successfully applied to the National Science Foundation, which approved \$18,500. ACL itself committed an additional \$10,000 in funding. Together, these sources provide approximately \$28,500 to directly support travel and participation for 28 student researchers. These combined funding efforts reflect the strong commitment of the ACL community to fostering early stage research and supporting equitable participation in the field.

We are deeply grateful to our faculty advisors, Yves Scherrer, Elias Stengel-Eskin, and Snigdha Chaturvedi, for their guidance, encouragement, and continued support. We thank all 374 members of the program committee for their thoughtful and thorough reviews.

We would like to recognize the mentors who contributed their time and expertise to guide our student authors: Mohamed Abdalla, Somak Aditya, Ameeta Agrawal, Georgios Alexandridis, Valerio Basile, Shubham Chatterjee, Fuxiang Chen, Haihua Chen, Hanjie Chen, Wei Chen, Ekapol Chuangsuwanich, Viviana Cotik, Dipankar Das, Sourish Dasgupta, Zhijie Deng, Kaize Ding, Khoa D Doan, Ondrej Dusek, Kripabandhu Ghosh, Tejas Gokhale, Venkata S Govindarajan, Nidhi Goyal, Sharath Chandra Guntuku, Qing Guo, Vivek Gupta, Mohammad Rashedul Hasan, Daniel Hershcovich, Xuming Hu, Kaiyu Huang, Oana Ignat, Labiba Jahan, Hyeju Jang, Yangfeng Ji, Lei Jiang, Zhuoxuan Jiang, Mandy Barrett Korpusik, Marek Kubis, Sebastian Kula, Sachin Kumar, Yen-Ling Kuo, Salem Lahlou, Allison Lahnala, Yuxuan Lai, Dung D. Le, Kevin Leach, Martha Lewis, Dongyuan Li, Sheng Li, Constantine Lignos, Gilbert Lim,

Lu Lin, Ying-Jia Lin, Dugang Liu, Yang Liu, Sharid Loáiciga, Dongsheng Luo, Chunchuan Lyu, Lorenzo Malandri, Chandresh Kumar Maurya, Fanchao Meng, Pruthwik Mishra, Raha Moraffah, Lili Mou, Rajesh Kumar Mundotiya, Hamada Nayel, Luyao Niu, Kunwoo Park, Jiahuan Pei, Mengyang Qiu, Leonardo Ranaldi, Yogesh S Rawat, Anthony Rios, Christophe Rodrigues, Jonathan Rusert, Yusuke Sakai, Manuela Sanguinetti, Debarshi Kumar Sanyal, Sunil Saumya, Salim Sazed, Lanyu Shang, Raksha Sharma, Wenqi Shi, A.B. Siddique, Gagandeep Singh, Mayank Singh, Nikhil Singh, Mohit Singhal, Konstantinos Skianis, Rui Sousa-Silva, Andreas Spitz, Ruixiang Tang, Yapeng Tian, Yu Tian, Cagri Toraman, Chenguang Wang, Dongxia Wang, Ruibo Wang, Yile Wang, Zach Wood-Doughty, Anna Wróblewska, Aiping Xiong, Depeng Xu, Han Xu, Lei Xu, Qionghai Xu, Wei Xue, Wenpeng Yin, Yu Yin, Chenyu You, Annie Zaenen, Chi Zhang, Qing Zhang, Xuhong Zhang, Linchao Zhu

Finally, we thank all authors for their enthusiasm, effort, and contributions. Your work makes the Student Research Workshop a vibrant and intellectually engaging part of ACL 2026.

We hope you find this year's workshop inspiring, welcoming, and enriching.

Juan Diego Rodriguez, Santosh T.Y.S.S., and Ona de Gibert  
ACL 2026 SRW Student Chairs

# Organizing Committee

## **Student Research Workshop Chairs**

Santosh T.Y.S.S., Technical University of Munich & Amazon  
Juan Diego Rodriguez, University of Texas at Austin  
Ona de Gibert, University of Helsinki

## **Student Research Workshop Chairs - Faculty Advisors**

Yves Scherrer, University of Oslo  
Elias Stengel-Eskin, University of Texas at Austin  
Snigdha Chaturvedi, University of North Carolina

# Program Committee

## Program Committee

Giuseppe Abrami, Johann Wolfgang Goethe Universität Frankfurt am Main  
Somak Aditya, Indian Institute of Technology Kharagpur  
Shardul Aggarwal  
Ameeta Agrawal, Portland State University  
Sina Ahmadi, University of Zurich  
Xi Ai  
Taichi Aida, Tokyo Metropolitan University, NEC and Hitotsubashi University  
Felermino D. M. A. Ali  
Pegah Alipoormolabashi, State University of New York at Stony Brook  
Diego Alves, Universität des Saarlandes  
Zaid Alyafeai, King Abdullah University of Science and Technology  
Evelin Amorim, Inesc tec  
Na Min An, NAVER and KAIST  
Xiang Ao, Institute of Computing Technology, Chinese Academy of Sciences and University of the Chinese Academy of Sciences  
Mario Ezra Aragon, Universidad de Santiago de Compostela  
Seth Aycock  
Long Bai, Institute of Computing Technology, Chinese Academy of Sciences  
JinYeong Bak, SungKyunKwan University  
Bal Krishna Bal, Kathmandu University  
Nishant Balepur  
Sivaji Bandyopadhyay  
Rachel Bawden, Inria  
Amanda Bertsch, Carnegie Mellon University  
Gagan Bhatia  
Arnab Bhattacharya, IIT Kanpur  
Johanna Björklund, Umea University  
Emanuela Boros, University of La Rochelle  
Daniel Braun, University of Marburg  
Khac-Hoai Nam Bui, Viettel Group  
Paul Buitelaar, University of Galway  
Rémi Cardon, Universidad Carlos III de Madrid  
Silvia Casola  
Dumitru-Clementin Cercel, National University of Science and Technology POLITEHNICA Bucharest  
Dong-Kyu Chae, Hanyang University  
Andong Chen  
Qiguang Chen  
Shuguang Chen, Purdue University  
Sishuo Chen, Alibaba Group  
Wei-Fan Chen, AIST, National Institute of Advanced Industrial Science and Technology  
Yidong Chen, Xiamen University  
Yulin Chen  
Wei Cheng, Nanjing University  
Xin Cheng, Peking University  
David Chiang, University of Notre Dame

Juhwan Choi, AITRICS  
Seungtaek Choi, Hankuk University of Foreign Studies  
Maximin Coavoux, CNRS  
Parag Pravin Dakle, Fidelity Investments  
Daniel Dakota, Leidos and Indiana University  
David Dale, FAIR at Meta  
Souvik Das, J.P. Morgan Chase  
Gaël de Chalendar, CEA, Paris-Saclay University  
Naihao Deng  
Sourabh Deoghare  
Jana Diesner, Technische Universität München  
Liang Ding, Alibaba Group  
Bin Dong, Ricoh Software Research Center Beijing Co., Ltd.  
Xiangjue Dong, Texas A&M University - College Station  
Abteen Ebrahimi, University of Colorado, Boulder  
Hafsteinn Einarsson, University of Iceland and deCODE genetics  
Roxanne El Baff, German Aerospace Center and Bauhaus-University Weimar  
Ignatius Ezeani, Lancaster University  
Martin Fajčík, Brno University of Technology  
Mariia Fedorova  
Aloka Fernando, University of Moratuwa  
Olivier Ferret, CEA-List  
Simona Frenda, Heriot-Watt University  
Yingxue Fu, INRIA  
Baban Gain, Indian Institute of Technology, Patna  
Yujian Gan, The Queen's University Belfast  
Hui Gao  
Jun Gao  
Shengxiang Gao, Kunming University of Science and Technology  
Michael Ginn, University of Colorado at Boulder  
Isao Goto, Ehime University  
Christan Grant, University of Florida  
Cyril Grouin, CNRS  
Jinghang Gu, Hong Kong Polytechnic University  
Tianle Gu  
Frank Guerin  
Gaël Guibon, University Paris 13, Université Paris Nord (Paris XIII)  
Ivan Habernal, Ruhr-Universität Bochum  
Patrick Haller, Humboldt Universität Berlin  
Hojae Han, Electronics and Telecommunications Research Institute  
Meng Han, Zhejiang University  
Annette Hautli-Janisz, Universität Passau  
Yoshihiko Hayashi, Waseda University  
Ruifang He  
Tosho Hirasawa, LayerX  
Carolin Holtermann, Universität Hamburg  
Christopher M. Homan, Rochester Institute of Technology  
Eftekhar Hossain, University of Central Florida  
David M Howcroft, University of Aberdeen  
Guimin Hu, Guangdong University of Technology  
Yiran HU

Frederikus Hudi, Nara Institute of Science and Technology, Japan  
Kathy Hämmerl  
Nikolai Ilinykh, Göteborg University  
Dmitry Ilvovsky, Higher School of Economics  
Mert Inan, Northeastern University  
Naoya Inoue, Ichikara, RIKEN and Japan Advanced Institute of Science and Technology  
Hasan Iqbal, Mohamed bin Zayed University of Artificial Intelligence  
Robert L. Logan IV, Dataminr  
Vladimir V. Ivanov  
Kokil Jaidka, National University of Singapore  
Dongjun Jang, Seoul National University  
Eugene Jang, Northeastern University  
Hyewon Jang, Göteborg University  
Inigo Jauregi Unanue, University of Technology Sydney  
Yuxiang Jia, Zhengzhou University  
Jiyue Jiang, The Chinese University of Hong Kong  
Shan Jiang, ebay Inc.  
Zhuoxuan Jiang, Shanghai Business School  
Kyohoon Jin, DATUMO  
Aditya Joshi, University of New South Wales  
Simeon Junker, Universität Bielefeld  
Tomoyuki Kajiwara, Ehime University  
Hung-Yu Kao, Department of Computer Science, National Tsing Hua University, National Tsing-hua University  
Manav Nitin Kapadnis, Apple  
Amir Hossein Kargaran  
Yoshihide Kato, Nagoya University  
Simerjot Kaur, JPMorgan Chase and Co  
Gyuwan Kim, University of California, Santa Barbara  
YoungBin Kim, Chung-Ang University  
Hirokazu Kiyomaru, National Institute of Informatics  
Rebecca Knowles, National Research Council Canada  
Philipp Koehn, Johns Hopkins University  
Mamoru Komachi, Hitotsubashi University  
Rik Koncel-Kedziorski, Apple  
Marek Kubis, Adam Mickiewicz University of Poznan  
Adithya Kulkarni, Ball State University  
Sonal Kumar  
Sujay S Kumar  
Prasanna Kumar Kumaresan, Data Science Institution, University of Galway, Ireland  
Jenny Kunz, Linköping University  
Yanis Labrak  
Tiziano Labruna  
Stefan Larson, Vanderbilt University  
Alberto Lavelli, Fondazione Bruno Kessler  
Seolhwa Lee  
Lauren Levine, Georgetown University  
Bobo Li, National University of Singapore  
Bryan Li, Google  
Chong Li, Nanyang Technological University  
Dongyuan Li, The University of Tokyo and Institute of Science Tokyo

Hao Li  
Jiatong Li, Hong Kong Polytechnic University  
Jiazheng Li, King's College London, University of London  
Junlin Li  
Lei Li, Beijing University of Posts and Telecommunications  
Mingzhe Li  
Ping Li, Southwest Petroleum University  
Shichen Li  
Yanzeng Li, Beijing Normal University  
Ying Li  
Yinghui Li, Tencent  
Yingya Li, Harvard University  
Zhuang Li, Royal Melbourne Institute of Technology  
Zihao Li  
Xiaobo Liang, Soochow University  
Pin-Jie Lin  
Chen Cecilia Liu, University of Cambridge  
Danni Liu, Karlsruher Institut für Technologie  
Dongqi Liu  
Jiacheng Liu  
Tianyu Liu, ETHZ - ETH Zurich  
Xiang Liu  
Yuanxing Liu, Harbin Institute of Technology  
Zheyuan Liu, University of Notre Dame  
Tien-Hong Lo, National Taiwan Normal University  
Yunfei Long, Queen Mary, University of London  
Natalia V Loukachevitch, Lomonosov Moscow State University  
Haoran Luo, Nanyang Technological University  
Saad Mahamood, Shopware  
Bhavivyva Malik, Edinburgh University, University of Edinburgh  
Hieu Man, University of Oregon  
Ramesh Manuvinakurike, Zillow Group, Inc.  
Marion Di Marco, Technische Universität München  
Magdalena Markowska  
Michele Mastromattei, Istituto Clinico Humanitas IRCCS - Humanitas Research Hospital  
Yuji Matsumoto, Tohoku University  
John Philip McCrae, National University of Ireland Galway  
Alexander Mehler, Johann Wolfgang Goethe Universität Frankfurt am Main  
Stephen Meisenbacher  
Alessio Miaschi, Institute for Computational Linguistics A. Zampolli"(CNR-ILC), Pisa  
Timothee Mickus, University of Helsinki  
Benjamin Minixhofer, Allen Institute for Artificial Intelligence and University of Cambridge  
Hideya Mino, NHK  
Pruthwik Mishra, Sardar Vallabhbhai National Institute of Technology  
Taro Miyazaki, NHK Science and Technology Research Laboratories  
Diego Molla, Macquarie University  
Ishani Mondal  
Lili Mou, University of Alberta  
Philippe Muller, IRIT, University of Toulouse  
Saliha Muradoglu  
Hongbin Na, University of Technology Sydney

Masaaki Nagata, NTT, Inc.  
Sudip Kumar Naskar  
Mariana Neves, German Federal Institute for Risk Assessment  
Vincent Ng  
Duc-Vu Nguyen, University of Information Technology, Vietnam National University  
Le-Minh Nguyen, Japan Advanced Institute of Science and Technology, Tokyo Institute of Technology  
Long HB Nguyen, Ho Chi Minh city University of Science, Vietnam National University  
Yuxiang Nie, Hong Kong University of Science and Technology  
Malvina Nikandrou  
Dmitry Nikolaev, University of Manchester  
Takashi Ninomiya, Ehime University  
Noriki Nishida, AIST, National Institute of Advanced Industrial Science and Technology  
Adam Nohejl, RIKEN  
Damien Nouvel, National Institute for Oriental Languages and Civilizations  
Odunayo Ogundepo  
Shu Okabe, Technische Universität München  
Maite Oronoz  
Naoki Otani, Megagon Labs  
Siqi Ouyang, CMU, Carnegie Mellon University  
Santanu Pal, Wipro  
Alexander Panchenko, S-NLP group  
Subhadarshi Panda, Samsung  
Liang Pang, Institute of Computing Technology, Chinese Academy of Sciences  
Denis Paperno, Utrecht University and CNRS  
Eunhwan Park, Motif Technologies  
Hyuntae Park, Korea University  
Kunwoo Park, Soongsil University  
Shinwoo Park  
Charith Peris, Amazon  
Irina Piontkovskaya, Huawei Technologies Ltd.  
Lidia Pivovarova, University of Helsinki  
Paul Piwek, Open University  
Esther Ploeger, Utrecht University  
Max Ploner, Humboldt Universität Berlin  
Giovanni Puccetti, CNR  
Sukannya Purkayastha, Bayerische Julius-Maximilians-Universität Würzburg  
Matthew Purver, Queen Mary University of London and Institut Jožef Stefan  
Leonardo Ranaldi  
Tharindu Ranasinghe, Lancaster University  
Sebastian Reimann, Ruhr-Universität Bochum  
Ehud Reiter, University of Aberdeen  
Nils Reiter, Universität zu Köln  
Myrthe Reuver, Populytics and Delft University of Technology  
MohammadHossein Rezaei, Scale AI  
Matīss Riktērs, National Institute of Advanced Industrial Science and Technology (AIST)  
Kyeongmin Rim  
Rudolf Rosa, Charles University, Prague  
Terry Ruas, Georg-August Universität Göttingen  
Jonathan Rusert, Purdue University Fort Wayne  
Yusuke Sakai, Nara Institute of Science and Technology, Japan

Tanja Samardzic, University of Zurich and IDSIA USI-SUPSI  
Fahime Same, Trivago N.V.  
Yves Scherrer, University of Oslo  
Patrícia Schmidtová  
Sabine Schulte im Walde, University of Stuttgart  
Shubhashis Sengupta  
Christophe Servan, AMIAD  
Matthew Shardlow, The Manchester Metropolitan University  
Sumit Shekhar, Adobe Systems  
Tianhao Shen  
Lei Shi, Örebro University  
Ning Shi, University of Alberta  
Anthony Sicilia, West Virginia University  
Lucia Siciliani, University of Bari  
Anil Kumar Singh, Indian Institute of Technology (Banaras Hindu University) Varanasi  
Mukul Singh, Microsoft  
Aman Sinha  
Xingyi Song, University of Sheffield  
Yixiao Song, University of Massachusetts at Amherst  
Alexey Sorokin, Lomonosov Moscow State University , Yandex and Lomonosov Moscow State University  
Saurabh Srivastava  
Sebastian Steindl, Ostbayerische Technische Hochschule Amberg-Weiden  
Elias Stengel-Eskin, University of Texas at Austin  
Ivelina Stoyanova, Deaf Studies Institute and Institute for Bulgarian Language, Bulgarian Academy of Sciences  
Sara Stymne, Uppsala University  
Sebastian Stüker, Zoom Video Communications  
Katsuhito Sudoh, Nara Women's University  
Mingming Sun, Beijing Institute of Mathematical Sciences and Applications  
Qiushi Sun, University of Hong Kong  
Anirudh Sundar, Microsoft  
Hanna Suominen, Australian National University  
Manan Suri  
Víctor M. Sánchez-Cartagena, Universidad de Alicante  
Fei Tan, East China Normal University  
Minghuan Tan, Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences, Chinese Academy of Sciences  
Jing Tang  
Philippe Thomas, German Research Center for AI  
Yu Tian  
Ryoko Tokuhisa, Aichi Institute of Technology  
Cagri Toraman, METU, Middle East Technical University  
David Traum, University of Southern California  
George Tsatsaronis  
Yu-Hsiang Tseng, Eberhard-Karls-Universität Tübingen  
Can Udomcharoenchaikit, Vidyasirimedhi Institute of Science and Technology  
Gökçe Uludoğan  
Masao Utiyama, National Institute of Information and Communications Technology (NICT), National Institute of Advanced Industrial Science and Technology  
Takehito Utsuro, University of Tsukuba

Teemu Vahtola, University of Helsinki  
Sowmya Vajjala, National Research Council Canada  
Soroush Vosoughi, Dartmouth College  
Bichen Wang, Harbin Institute of Technology  
Chenglong Wang, Northeastern University  
Chengyu Wang, Independent Researcher and Alibaba Group  
Daling Wang  
Jiapeng Wang  
Qingyun Wang, College of William and Mary  
Shuai Wang, Wuhan University  
Siyin Wang, Department of EE, Tsinghua University  
Song Wang, University of Central Florida  
Yimu Wang, University of Waterloo  
Zimu Wang, Monash University and University of Liverpool  
Leo Wanner, Barcelona Supercomputing Center and Catalan Institute for Research and Advanced Studies (ICREA)  
Bonnie Webber, Edinburgh University, University of Edinburgh  
Bosi Wen  
Steven R Wilson, University of Michigan - Flint  
Genta Indra Winata, Capital One  
Christian Wolff, University of Regensburg  
Zach Wood-Doughty, Northwestern University  
Fangzhao Wu, Microsoft  
Zekun Wu, Department of Computer Science, University College London, University of London and Holistic AI  
Heming Xia, Hong Kong Polytechnic University  
Ziang Xiao, Department of Computer Science, Whiting School of Engineering  
Kaige Xie, Georgia Institute of Technology  
Shufang Xie, Renmin University of China  
Bing Xu  
Chen Xu, Harbin Engineering University  
Han Xu, University of Arizona  
Shanshan Xu, University of Copenhagen  
Tianyang Xu, Purdue University  
Wei Xue  
Kosuke Yamada, CyberAgent, Inc.  
Atsuki Yamaguchi, University of Sheffield  
Ivory Yang, Dartmouth College  
Li Yang, Institute of Software, Chinese Academy of Sciences  
Tianyu Yang  
Lina Yao, University of New South Wales  
Qinyuan Ye, Salesforce Research  
Reyyan Yeniterzi, GenAIus Technologies  
Yuwei Yin, University of British Columbia  
Haneul Yoo, New York University  
Eunseop Yoon, KAIST  
Koichiro Yoshino, Osaka University and RIKEN  
Bei Yu, Department of Computer Science and Engineering, The Chinese University of Hong Kong  
Hao Yu, National University of Defense Technology  
Linhao Yu  
Seunguk Yu, Chung-Ang University

Zhengtao Yu, Kunming University of Science and Technology and Yunnan Key Laboratory of Artificial Intelligence  
Caixia Yuan  
Ye Yuan  
JungMin Yun, Chung-Ang University  
Andrea Zaninello  
Qingkai Zeng, Nankai University  
Runzhe Zhan, University of Macau  
Chen Zhang, Peking University  
Feng Zhang, Peking University  
Haoyu Zhang, The Chinese University of Hong Kong, Shenzhen and Tencent LightSpeed Studios  
Hongyu Zhang, Chongqing University  
Jian Zhang  
Licheng Zhang  
Mike Zhang, University of Copenhagen  
Ningyu Zhang, Zhejiang University  
Xuanming Zhang, Columbia University  
Yi Zhang, University of Zurich and ZHAW - Zürcher Hochschule für Angewandte Wissenschaften  
Yichi Zhang, Zhejiang University  
Zhong Zhang, University of Electronic Science and Technology of China  
Jie Zhao, Chang'an University  
Xiangyu Zhao, City University of Hong Kong  
Xingmeng Zhao, University of Colorado Anschutz Medical Campus  
Xiaoqing Zheng, Fudan University  
Yang Zhong, University of Pittsburgh  
Baohang Zhou, Tiangong University  
Ben Zhou, Arizona State University  
Fan Zhou, University of Electronic Science and Technology of China  
Henghui Zhu, Figma  
Junnan Zhu, Institute of automation, Chinese academy of science, Chinese Academy of Sciences  
Tong Zhu, Shanghai Artificial Intelligence Laboratory  
Xinlin Zhuang, Shopee  
Maike Züfle, Karlsruher Institut für Technologie

## Table of Contents

<i>Multi-Agent Reasoning Improves Compute Efficiency: Pareto-Optimal Test-Time Scaling</i> Florian Valentin Wunderlich, Lars Benedikt Kaesberg, Jan Philip Wahle, Terry Ruas and Bela Gipp .....	1
<i>Claim Verification in the Age of Large Language Models: A Survey</i> Alphaeus Dmonte, Roland R Oruche, Marcos Zampieri, Prasad Calyam and Isabelle Augenstein	15
<i>Language Directions in Multilingual LLMs: A Layer-wise Diagnostic Study of Token Alignment and Pretraining Imprint</i> JaeSeong Kim and Suan Lee .....	30
<i>Emergence of Minimal Circuits for Indirect Object Identification in Attention-Only Transformers</i> Rabin Adhikari .....	36
<i>Why Large Language Models can Secretly Outperform Embedding Similarity in Information Retrieval</i> Matei Benescu and Ivo Pascal de Jong .....	47
<i>Detecting Hallucinations in Large Language Models via Internal Attention Divergence Signals</i> Gijs van Dijk .....	60
<i>Reflection in the Dark: Exposing and Escaping the Black Box in Reflective Prompt Optimization</i> Shiyun Liu, Qifeng Xia, Qiyun Xia, Yisheng Liu, Xinyu Yu and Rui Qu .....	76
<i>Thesis Proposal: LLMs post-training for multilingual medical tasks. Instruction-Tuning, Continual-Pretraining or Reasoning?</i> Pietro Ferrazzi, Alberto Lavelli and Bernardo Magnini .....	110
<i>Peek2: Regex-free Byte-level Byte-Pair Encoding Pretokenizer for LLM Inference on Edge Devices</i> Liu Zai and Iraklis A. Klampanos .....	123
<i>Annotation Entropy Predicts Per-Example Learning Dynamics in LoRA Fine-Tuning</i> Brady Steele .....	129
<i>Semantic Contrastive Adaptation for Multimodal Figurative Language Understanding</i> Ayaan Siddiqui .....	142
<i>Think Less, Code Better: Probing When Chain-of-Thought Hurts and How to Route Around It</i> Rajarshi Ghoshal, Debadri Basak, Salma Emad Mahmoud Abdelhalim and Pratibha Kaur Arora	152
<i>Probing Functional Correctness in Diffusion Language Models</i> Guan-Ming Chiu and Jeng-Yue Liu .....	163
<i>Thesis Proposal: On the Granularity-Robustness Trade-off in Text-Derived Knowledge Graphs</i> Surawat Pralomram .....	173
<i>TokLens: A Multilingual Lens on Tokenizer Quality for LLMs</i> Guan-Ming Chiu .....	188
<i>Phase Transitions in Affective Meaning Divergence: The Hidden Drift Before the Break</i> Napassorn Litchiowong .....	206

<i>Sycophantic Anchors: Localizing and Quantifying User Agreement in Reasoning Models</i>	
Jacek Duszenko, Przemyslaw Kazienko and Jan Kocon .....	225
<i>NEAT-IR: Neural Explainable Analysis Tool for Information Retrieval</i>	
Lev Sukherman, Artem Frenk, Nina Klimenkova and Connor Jason .....	240
<i>BanglaSocialBench: A Benchmark for Evaluating Sociopragmatic and Cultural Alignment of LLMs in Bangladeshi Social Interaction</i>	
Tanvir Ahmed Sijan, S. M. Golam Rifat, Pankaj Partha, Md. Tanjeed Islam and Md Musfique Anwar .....	247
<i>Interpretability of LLM Classifiers via the Rational Inattention Theory with Application to Hate Speech Detection</i>	
Yuan Zhao and Ali Abdi .....	281
<i>The Shape of Vulnerability: How Adversarial Perturbations Reshape the Topology of Language Model Latent Spaces</i>	
Angelina Tsai, Shreya Subramanian, Catherine Liu, Kimberly Lopez, Leif Zinn-Brooks, Alexia E. Schulz and Adaku Uchendu .....	290
<i>LLM-based Literal Example Generation for Japanese Multiword Expressions</i>	
Mio Ohashi, Hajime Kiyama, Zhidong Ling and Mamoru Komachi .....	307
<i>Presentation Slide Translation and Layout Error Correction by LLMs</i>	
Futo Kajita, Nobuyori Nishimura, Takehito Utsuro, Naoki Muto, Chee Siang Leow and Hiromitsu Nishizaki .....	326
<i>Constructing a Japanese Rap Lyric Generation Model with GRPO</i>	
Hayato Ogawa and Daisuke Kawahara .....	340
<i>Tracking the Evolution of Foresight Signals in News Data: The Case of the European Electric Vehicle Market</i>	
Karine Navasartian .....	352
<i>Debiasing Logical Fallacy Detection for Real-World Robustness via Counterfactually Augmented Data</i>	
Navyansh Singh .....	363
<i>Bring Linguistics Back to Cryptanalysis: Using Attestation to Break the Advanced Encryption Standard</i>	
Madeline Boese .....	375
<i>Garden Path Recovery in Causal and Masked Language Models</i>	
Sanjan Baitalik and Rajashik Datta .....	383
<i>Confidence as a Tie-Breaker: Reassessing Multilingual Hedging Bias in LLM-as-a-Judge Evaluation</i>	
Rajashik Datta and Sanjan Baitalik .....	393
<i>BanglaSTEM: A Parallel Corpus and Term-Weighted Evaluation for Technical Bangla-English Translation</i>	
Kazi Reyazul Hasan, A. B. M. Alim Al Islam and Muhammad Abdullah Adnan .....	403
<i>Disentangling the Effects of Unlearning in Measuring Parametric Faithfulness of Chain-of-Thought</i>	
Ryo Mitsuhashi, Gaku Morio, Ayana Niwa, Masahiro Kaneko, Kentaro Inui, Terufumi Morishita, Yuta Koreeda and Yasuhiro Sogawa .....	413
<i>FedPAGR: Federated Prototype Alignment via Geometric Refinement for Heterogeneous Architectures</i>	
Kris Prasad and Md Abdullah Al Hafiz Khan .....	420

<i>Sentiment Analysis of Yelp Review Dataset: A Comparative Study of Machine Learning Methods</i> Krishna Thakar, Mohamed Abu Sheha and Emmanuel Thompson .....	428
<i>Semantic Span Annotation: An Exploratory Study of LLM Annotation</i> Tejas Goyal, Dhriti Krishnan, Anuj Gupta and Jaromir Savelka .....	438
<i>Thesis Proposal: An Explainable Multimodal Framework for Detecting Harmful Content in Code-Switched Children’s Media</i> Juliana Isabelle A. Guillermo, Jasper Kyle Catapang and Nathaniel Oco .....	450
<i>Test-Time Strategies for More Efficient and Accurate Agentic RAG</i> Abhinav Sharma, Brian Zhang, Deepti Guntur, Zhiyang Zuo, Shreyas Chaudhari, Wenlong Zhao, Franck Dernoncourt, Puneet Mathur, Ryan A. Rossi and Nedim Lipka .....	463
<i>Eye Movement Features Can Predict Human Preferences on Machine-Generated Texts</i> Xiaoshan He, Xiaoqun Liu, Haodong He, Yu Wang and Yang Xu .....	470
<i>Thesis Proposal: Diagnosing and Mitigating Semantic Interference in Script-Sharing Low-Resource Language Models: A Case Study on Square Bai Script</i> Jingting Zheng and Deyi Xiong .....	487
<i>Does Locality Cost in Polish Medical Text Classification? Duplicate-Aware Evaluation of Federated Learning</i> Daniel Cieślak and Andrzej Czyżewski .....	498
<i>Analyzing Hate Speech Amplification on Fringe Platforms</i> Anika Ghosh Basu, Humberto Jesus Carlon and Junyi Liu .....	509
<i>The Silence of the Facts: Popularity as a Barrier to Machine Unlearning</i> Anna Borisiuk, Andrey Savchenko, Alexander Panchenko and Elena Tutubalina .....	515
<i>Leakage-Aware User-Level ADHD Signal Classification from Social Media: When Graph Aggregation Helps, and When It Does Not</i> Daniel Cieślak and Władysław Średniawa .....	527
<i>CAL-Log: Cost-Aware Active Learning with Logarithmic Cognitive Effort Modeling and Online Adaptation to Human Annotation Behavior</i> Vihanga Supasan Kariyakaranage and Banuka Athuraliya .....	537
<i>Thesis Proposal: Targeted and Unified Cross-Lingual Unlearning from Multilingual Language Models</i> Jan Bronec and Jindřich Helcl .....	554
<i>AllY-Compressor: A Framework for Enhancing the Efficiency of GUI Agent Observations through Visual Context Reconstruction and Redundancy Reduction</i> Michito Takeshita, Takuro Kawada, Takumi Ohashi, Shunsuke Kitada and Hitoshi Iyatomi ..	563
<i>Counterspeech Generation using Small Language Models</i> Abubakar Sadiq Muhammad, Simona Frenda and Gavin Abercrombie .....	581
<i>From Graphs to Hypergraphs: Enhancing Aspect-Term Sentiment Analysis via Multi-Level Relational Modeling</i> Omkar Mahesh Kashyap, Padegal Amit, Madhav Kashyap, Ashwini M Joshi and Shylaja S S	595
<i>Probing Bias Formation in Medical LLMs through Activation Steering</i> Bayram Ayadi and Annette Hautli-Janisz .....	609
<i>Faithfulness Beyond Plausibility: Auditing Human Explanations in Educational Assessment</i> Ria Talsania, Dhruv Ritesh Shah and Sudhir Dhage .....	621

<i>CBAL: Context-Based Agentic Learning for Speaker Diarization Segmentation Refinement</i> Odwitiyo Dutta and Dinesh K Vishwakarma .....	632
<i>Measuring and Mitigating Shortcut Reliance in Language Models with Probe-Based Representation Entanglement</i> Divyajot Singh .....	648
<i>LAMP-MedQA: A Lightweight Multi-Agent System for Patient-Oriented Medical Question Answering</i> Jack A. Johnson, Meghali Banerjee, Joseph Crawford, James Welch, Jim Davies and Tingyan Wang .....	663
<i>Inference-Time Feedback for Reasoning Controllability in Diffusion Language Models</i> Clovis Barbour and Huixin Zhan .....	675
<i>Disentangling Linguistic Relatedness from Task Alignment in Cross-Lingual Transfer</i> Ahmed Haj Ahmed, Ruochen Zhang and Alvin C Grissom II .....	683
<i>PE-QAT: Parameter-Efficient Quantization-Aware Training for Large Language Models</i> Shresth Mishra .....	701
<i>Fusion Training for Mathematical Generalization in Large Language Models</i> Congfeng Cao, Pengyu Zhang and Jelke Bloem .....	712
<i>Does Topic Sentiment Cause Perceived Ideology? Comparing Human and LLM Annotations in Political News Articles</i> Upasana Chatterjee .....	725
<i>Understanding Conversational Implicatures in Humans and LLMs</i> Daeun Kang .....	741
<i>Thesis Proposal: Self-Adaptive and Epistemic Uncertainty-Guided ASR of Dense Intra-Sentential Code-Switched Speech for African Low-Resource Languages</i> Umar Baba Umar, Sulaimon Adebayo Bashir, Abdulmalik Danlami Mohammed and Amina Gogo Tafida .....	754
<i>RegTrack: A Fine-Grained Benchmark for Multi-Class Legal Change Detection</i> Joe Yu, Kevin Chenhao Li and Julian Ostarek .....	764
<i>Validator-Guided Hard Negative Mining for Masked Language Modeling in Low-Resource Ancient Languages</i> Andrei Voinea .....	779
<i>Conformal LLM Routing with Distribution-Free Safety Guarantees</i> Iqtedar Uddin and André Bauer .....	791
<i>Supervision versus Demonstration-Based In-Context Learning for Multiword Expression Classification</i> Sercan Karakaş and Yusuf Şimşek .....	800
<i>When Models Hesitate: Answer Instability as a Label-Free Uncertainty Signal for LLMs</i> Jasper Meynard Arana, Kristine Ann M. Carandang, Ethan Robert Casin, Christian Alis and Christopher Monterola .....	816
<i>Thesis Proposal: Toward a Human-Centered and Perspective-Aware Framework for Reproducible ML Evaluation and AI Alignment</i> Deepak Pandita and Christopher M. Homan .....	827
<i>One Panel Does Not Fit All: Case-Adaptive Multi-Agent Deliberation for Clinical Prediction</i> Yuxing Lu, Yushuhong Lin and Jason Zhang .....	844

<i>LLMs for Now, Fine-Tuning for Later: An Ensemble Approach to Data Drift in Domain-Specific Tasks</i> Yuxuan Lu, Bingsheng Yao, Shao Zhang, Yisi Sang, Yun Wang, Hansu Gu, Peng Zhang, Tun Lu, Toby Jia-Jun Li and Dakuo Wang .....	861
<i>Thesis Proposal: When Does an Agent Know It Is Lost? Confidence Trajectory Analysis for Tool-Using LLMs</i> Zhenjiang Mao .....	877
<i>Task Assignment meets Annotator Modeling: Human-LLM Collaborative Annotation with Constraints</i> Kei Moriyama, Kouta Nakayama and Yukino Baba .....	888
<i>Dynamic Meta-Metrics: Source-Sentence Conditioned Weighting for MT Evaluation</i> Luke Zhang, Justin Vasselli, Aditya Khan, York Hay Ng and En-Shiun Annie Lee .....	903
<i>Tonal Saliency in Cognitive Decline: In-Context MCI Detection with Multimodal LLMs</i> Christopher Song and Abdullah P. Rashed Ahmed .....	913
<i>What Moves the Pareto Frontier in Tool-Using Agents? A Compute-Aware Study of ReAct Variants</i> Rishi N. Simhadri .....	923
<i>Filling the Long Tail: Structure-Aware Curriculum-Gap Completion for Medical Education with LLMs</i> Wenjie Lin .....	941
<i>Mechanistic Analysis Of Universality: Numerical Comparison Circuits Across Transformer Architectures</i> Arya Bhardia, Julian Ramirez, Siddhanta Verma and Karen Mkrtchyan .....	951
<i>How Hard is Math? Using Quantitative Metrics to Measure LLM Alignment to Human Intuitions of Difficulty</i> Micah Helzerman, Steven R Wilson and Cam McLeman .....	968
<i>Fine-Grained Semantic Comparison of Legal Documents using LLMs</i> Elisei Rykov, Nikolay Ivanov, Maria Bandulevich, Kseniia Petrushina, Valentin Malykh, Vasily Konovalov, Alexander Panchenko and Ilseyar Alimova .....	982
<i>From Fluent to Useful: Generative AI That Models Purpose, Audience, and Presenter for Scientific Communication</i> Ishani Mondal .....	995
<i>Beyond Discrete Search: Divergent Thinking as Intention Optimization in Latent Space</i> Mateusz Bystroński, Grzegorz Piotrowski and Tomasz Jan Kajdanowicz .....	1007
<i>Boosting Self-Consistency with Ranking</i> Maria Marina, Daniil Moskovskiy, Sergey Pletenev, Mikhail Salnikov, Alexander Panchenko and Viktor Moskvoretskii .....	1017
<i>LLM-Based Zero-Shot Soft Labeling for Anticipating Disagreement in Negotiation Dialogues</i> Ken Watanabe and Katsuhide Fujita .....	1033
<i>Analysis of the Neglect-Zero Effect in Large Language Models</i> Jin Tanaka, Daiki Matsuoka, Ryoma Kumon and Hitomi Yanaka .....	1052
<i>Morphology-Aware Multi-Granularity Representation Learning for Agglutinative Languages</i> Zhonghao Zhang, Na Liu, Jijia Ma, Nier Wu and Guiping Liu .....	1065
<i>Processing Inconsistency Predicts Language Competence: LLM Evaluation Without Answer Labels on Turkic Languages</i> Ilya Galyukshev and Ilseyar Alimova .....	1074

<i>TableMBR: Minimum Bayes Risk Table Generation Based on Structural Consistency</i>	
Daiki Yoshida, Hiroyuki Deguchi, Yusuke Sakai, Hidetaka Kamigaito and Taro Watanabe . . .	1087
<i>Through the Looking Glass of Multilingual AI: Contrasting Language- and Name Script-Dependent Ethnic Hierarchies in GPT and DeepSeek</i>	
Annabella Sakunkoo and Jonathan Sakunkoo . . . . .	1103
<i>Thesis Proposal: Auditing and Mitigating Demographic Bias in Multi-Stage Retrieval Systems for Criminal Justice Applications</i>	
Archan Dutta . . . . .	1115
<i>Contextual Diversity Measure (CDM) for Controllable Story Generation in Large Language Models</i>	
Richard Susilo, Hanna Suominen and Patrik Haslum . . . . .	1125
<i>Constructing a Japanese Verdict Prediction Dataset for Fact-Checking of LLM-Generated Texts</i>	
Miwa Masano, Hirokazu Kiyomaru, Atsushi Keyaki, Kaito Horio, Rei Minamoto, Ribeka Keyaki, Kouta Nakayama, Hideyuki Tachibana and Daisuke Kawahara . . . . .	1139
<i>Calibration vs Decision Making: Revisiting the Reliability Paradox in Unlearned Language Models</i>	
Divyaksh Shukla and Ashutosh Modi . . . . .	1152
<i>Disentangling Meaning and Language Components in Diverse Multilingual Sentence Embeddings</i>	
Kanade Nonomura, Keita Fukushima, Risa Kondo and Tomoyuki Kajiwara . . . . .	1169
<i>Linguistically-Informed Evaluation of LLMs on Acceptability Judgments in a Forced-Choice Paradigm</i>	
Ziyue Liu and Nils Reiter . . . . .	1177
<i>Representing Lean Proofs as Trajectories in Latent Space</i>	
Elisaveta Samoylov and Soroush Vosoughi . . . . .	1190
<i>Evaluation of Multilingual Ability to Use Spatial Deictic Expressions in Vision-Language Models</i>	
Kaito Watanabe, Taisei Yamamoto, Tomoki Doi and Hitomi Yanaka . . . . .	1203
<i>LLM Parameters for Math Across Languages: Shared or Separate?</i>	
Behzad Shomali, Luisa Victor, Tim Selbach, Ali Hamza Bashir, David Berghaus, Joachim Koehler, Mehdi Ali and Markus Frey . . . . .	1212
<i>Thesis Proposal: Intentional Inference for Insight Generation</i>	
Kristýna Onderková . . . . .	1236
<i>Reference-Free Schema Generation for Literature Review Tables via Multi-Faceted Rewards</i>	
Sinjoy Saha, Suman Saha, Mahfuza Farooque and Wenpeng Yin . . . . .	1253
<i>Factual State Discovery Benchmark: Evaluating Fact Elicitation in Polish Tax Law</i>	
Mateusz Bystroński, Kamil Tagowski, Denis Janiak, Julia Farganus, Lukasz Augustyniak, Monika Kajdanowicz and Tomasz Jan Kajdanowicz . . . . .	1262
<i>Evolutionary Search for Automated Design of Uncertainty Quantification Methods</i>	
Mikhail Seleznyov, Daniil Korbut, Viktor Moskvoretskii, Oleg Somov, Alexander Panchenko and Elena Tutubalina . . . . .	1280
<i>Thesis Proposal: A Normalization-First Framework for Sound, Complete, and Utility-Ready Open Information Extraction</i>	
Chandan Prakash, Pavan Kumar Chittimalli and Arnab Bhattacharya . . . . .	1291
<i>Mind the Gap: Multilingual Divide in LLM Bias Detection and Reasoning</i>	
Medha Hira, Prachi Goyal, Raj Maheshwari and Arnav Goel . . . . .	1305

<i>Multi-Constraint State Tracking with Negation: A Diagnostic Benchmark for LLM World Modeling</i> Ayan Sar, Pranav Singh Puri, Sumit Aich, Anurag Kaushish, Tanupriya Choudhury and Ajith Abraham .....	1317
<i>Learning Shortcut Models for Efficient Recursive Reasoning</i> Shiv Shankar .....	1351
<i>The Confident Liar: Diagnosing Multi-Agent Debate with Log-Probabilities and LLM-as-Judge</i> Ali Keramati, Justin Cheok, Jacob Horne and Mark Warschauer .....	1361
<i>Convergent Demographic Utility Hierarchies: Geometry of Intersectional Values in LLMs</i> Pravish Sainath .....	1376
<i>CRL-Prompt: Contrastive and Reinforcement Learning for Soft Prompt Tuning for Text Classification</i> Danila Lapokin and Andrey Savchenko .....	1391
<i>Optimizing Packing and Shuffling Strategies for Enhanced Performance in Generative Language Models</i> Yanbing Chen, Ruilin Wang, Zihao Yang, Lavender Yao Jiang and Eric Karl Oermann .....	1399
<i>LLM as a Meta-Judge: Synthetic Data for NLP Evaluation Metric Validation</i> Lukáš Eigler, Jindřich Libovický and David Hurych .....	1417
<i>Continuous Context Sampling Allows Extending Diversity Boundaries of Large Language Models</i> Mateusz Bystroński, Doheon Han, Nitesh V. Chawla and Tomasz Jan Kajdanowicz .....	1436
<i>One Task Vector is not Enough: A Large-Scale Study for In-Context Learning</i> Pavel Tikhonov, Ivan Oseledets and Elena Tutubalina .....	1451