

EACL 2026

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

Proceedings of the Student Research Workshop

March 26, 2026

©2026 Association for Computational Linguistics

Order copies of this and other ACL proceedings from:

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

ISBN 979-8-89176-383-8

Preface to the EACL 2026 Student Research Workshop

Welcome to the EACL 2026 Student Research Workshop (SRW)! This year marks a historic milestone as EACL is hosted in the African region for the very first time. We are thrilled to gather in a location that reflects the growing global footprint of our field, and we are eager to see the unique perspectives and diverse research directions this new setting brings to the workshop. As an integral part of the conference, the SRW provides students with a platform to present their work during the main conference and engage in meaningful dialogue with peers and experts within the NLP community.

This year, we received a record-breaking 185 submissions from students across the globe, which is a threefold increase compared to the most recent EACL 2024 SRW. This surge in interest is a clear testament to the vitality and expansion of NLP research. However, such rapid growth brought significant logistical challenges. To maintain our standards of quality while managing this volume, we had to rapidly scale our operations by recruiting a larger panel of 231 reviewers. To ensure expert feedback, we invited candidates from previous SRW events and the recent ACL Rolling Review (ARR) outstanding reviewer lists, requiring that all reviewers have at least one first-author publication in major NLP venues. Their dedication ensured that 94.5% of submissions received at least three reviews.

The selection process was highly competitive. Each submission underwent a rigorous double-blind review assessing its originality, technical quality, and clarity. From our record pool of 115 long papers, 54 short papers, and 16 thesis proposals, we accepted 77 papers in total (56 long, 12 short, and 9 thesis proposals). This resulted in an overall acceptance rate of approximately 41.6%. Of the accepted papers, 70 are archival and 7 are non-archival. We are proud of the high caliber of research represented in these proceedings.

Following the workshop's tradition, we also organized a pre-submission mentorship program to help student researchers refine their work before formal submission. Mirroring the surge in main submissions, the mentorship program saw a significant increase in participation with 53 submissions. We are incredibly grateful to the 27 mentors who volunteered their time. To ensure the highest level of guidance, we invited mentors with established track records, including those with a minimum of four publications in major NLP venues and a PhD in a relevant field.

The mentorship program is a crucial and unique feature of the SRW, contributing to our mission to foster growth and development in the NLP community. To further improve this experience, we conducted a survey among participants which provided valuable insights for future iterations. Notably, for the first time, we have included a mentorship program report in the proceedings. We hope this will serve as a valuable resource for the community and a roadmap for organizers of future mentorship programs.

Finally, an undertaking of this magnitude is only possible through the collective effort of a dedicated community. We would like to express our deepest gratitude to the General Chair, Aline Villavicencio; the Program Chairs, Vera Demberg, Kentaro Inui, and Lluís Marquez Villodre; and the ACL Business Manager, Jennifer Rachford, for their support and guidance throughout the process. We also thank the ARR workflow manager, Holy Lovenia, for her cooperation in detecting double submissions. We are equally grateful for the immense support from our Faculty Advisors, Tanja Samardzic and Marten van Schijndel, who provided invaluable advice and mentorship throughout the organization of the workshop. Finally, our thanks go to the reviewers and mentors for their tireless work, and to the authors for their contributions. We hope you find the workshop inspiring and that it sparks new collaborations within our global community.

Selene Báez Santamaría
Sai Ashish Somayajula
Atsuki Yamaguchi

EACL 2026 Student Research Workshop Chairs

Organizing Committee

Program Chairs

Selene Báez Santamaría, University of Zurich
Sai Ashish Somayajula, Oracle AI
Atsuki Yamaguchi, University of Sheffield

Faculty Advisors

Tanja Samardzic, IDSIA USI-SUPSI
Marten van Schijndel, Cornell University

Program Committee

Mentors

Omri Abend, Hebrew University of Jerusalem
Gavin Abercrombie, Heriot-Watt University
Zeljko Agic, Unity Technologies
Xiang Dai, CSIRO
Brian Davis, Dublin City University
Mattia Di Gangi, DeepL
Micha Elsner, Ohio State University
Neele Falk
Anmol Goel, Technische Universität Darmstadt
Hila Gonen, University of British Columbia
Ivan Habernal, Ruhr-Universität Bochum
Christopher Homan
Helmut Horacek
Amirhosein Javadi
Lea Krause, Vrije Universiteit Amsterdam
Marianne Liu, Oracle
Filip Miletic, University of Stuttgart
Amita Misra, Amazon
Ted Pedersen, University of Minnesota - Duluth
Daniel Peterson, Oracle Labs
Van-Thuy Phi, RIKEN
Joe Stacey, University of Sheffield
Xingwei Tan, University of Sheffield
Sowmya Vajjala, National Research Council Canada
David Vilar, Google
Bonnie Webber, University of Edinburgh
Yang Xu, University of Toronto
Olga Zamaraeva, Universidad de La Coruña

Reviewers

Gavin Abercrombie, Heriot-Watt University
Taichi Aida, Tokyo Metropolitan University, NEC and Hitotsubashi University
V.S.D.S.Mahesh Akavarapu, Eberhard-Karls-Universität Tübingen
Miguel Alonso, Universidade da Coruña
Evelin Amorim, Inesc tec
Amith Ananthram, Columbia University
Tatiana Anikina, German Research Center for AI
Mario Aragon, Universidad de Santiago de Compostela
Samee Arif
David Arps, HHU Düsseldorf
Yuwei Bao, Microsoft
Rachel Bawden, Inria
David Beauchemin
Tadesse Belay
María Benavente, Universidad Carlos III de Madrid

Himanshu Beniwal
Gabriel Bernier-Colborne, National Research Council Canada
Dario Bertero, Legality Ltd.
Manik Bhandari, Meta
Gagan Bhatia
Laura Biester, Middlebury College
Matteo Bortoletto
Davide Buscaldi, Ecole polytechnique and Université Paris 13
Elena Cabrio, Université Côte d'Azur
Samuel Cahyawijaya, Cohere
Mingzi Cao
Giovanni Cassani, Tilburg University
Jan Cegin, Kempelen Institute of Intelligent Technologies
Tanise Ceron
Ting-Yun Chang
Anwoy Chatterjee, Indian Institute of Technology, Delhi
Beiduo Chen, Ludwig-Maximilians-Universität München
Huiyao Chen
Ruirui Chen, Institute of High Performance Computing, Singapore, A*STAR
Andong Chen
Luis Chiruzzo, Facultad de Ingeniería - Universidad de la República - Uruguay
Elena Chistova
Juhwan Choi, AITRICS
Arijit Chowdhury, xAI and Amazon
Nico Daheim, ETHZ - ETH Zurich and Technische Universität Darmstadt
David Dale, FAIR at Meta
Naihao Deng
Jingcheng Deng
Sourabh Deoghare
Barbara Di Eugenio, University of Illinois, Chicago
Ritam Dutt, Carnegie Mellon University
Nils Dycke, TU Darmstadt
Esra Dönmez, Universität Stuttgart
Mohamed Elaraby
Micha Elsner, Ohio State University
Carlos Escolano, Barcelona Supercomputing Center
Neele Falk
Nils Feldhus
Thomas François, UCL
Giacomo Frisoni
Diana Galvan-Sosa
Yujian Gan, The Queen's University Belfast
Josef Genabith
Kshitish Ghate, University of Washington
Karim Ghonim, University of Roma La Sapienza"
Ameya Godbole, University of Southern California
Anmol Goel, Technische Universität Darmstadt
Sujatha Das Gollapalli, National University of Singapore
Dhiman Goswami, George Mason University
Nidhi Goyal, Mahindra University
Hojae Han, Electronics and Telecommunications Research Institute

Maram Hasanain, Qatar Computing Research Institute
Wei He
Shohei Higashiyama, Nara Institute of Science and Technology, Japan and National Institute of Information and Communications Technology (NICT), National Institute of Advanced Industrial Science and Technology
Tatsuya Hiraoka, Nara Institute of Science and Technology, Mohamed bin Zayed University of Artificial Intelligence and RIKEN
Christopher Homan
Helmut Horacek
Yuncheng Hua
Anthony Hughes
Filip Ilievski, Vrije Universiteit Amsterdam
Mert Inan, Northeastern University
Go Inoue, Mohamed bin Zayed University of Artificial Intelligence
Mete Ismayilzada, EPFL - EPF Lausanne
Jihyoung Jang, Pohang University of Science and Technology
Eugene Jang, Northeastern University
Fan Jiang
Pritam Kadasi, Indian Institute of Technology, Gandhinagar
Hour Kaing, National Institute of Information and Communications Technology (NICT), National Institute of Advanced Industrial Science and Technology
Tomoyuki Kajiwara, Ehime University
Pride Kavumba, SB Intuitions
Daisuke Kawahara, Waseda University
Amr Keleg, Mohamed bin Zayed University of Artificial Intelligence
Yova Kementchedjhiya, Mohamed bin Zayed University of Artificial Intelligence
Geewook Kim, University of Seoul, NAVER Cloud and KAIST
Brendan King, University of California, Santa Cruz
Philipp Koehn, Johns Hopkins University
Mamoru Komachi, Hitotsubashi University
Sanjeev Kumar, Indian Institute of Technology, Bombay
Jenny Kunz, Linköping University
Kemal Kurniawan, University of Melbourne
Matthieu Labeau, Télécom ParisTech
Allison Lahnala, McMaster University
Young-Jun Lee
Yuepei Li
Dongyuan Li, The University of Tokyo and Institute of Science Tokyo
Yinghui Li
Chao-Chun Liang, Asia University
Jasy Suet Yan Liew, Universiti Sains Malaysia
Constantine Lignos, Brandeis University
Gili Lior, The Hebrew University of Jerusalem
Zhu Liu
Danni Liu, Karlsruher Institut für Technologie
YongKang Liu, Northeast University
Chun Hei Lo, Huawei Technologies Ltd.
Wanqiu Long
Natalia Loukachevitch, Lomonosov Moscow State University
Agnes Luhtaru, University of Tartu
Stephanie Lukin, DEVCOM Army Research Laboratory

Ziyang Luo, Salesforce Research
Weiqing Luo, Arizona State University
Pedro Henrique Luz de Araujo, Universität Vienna
Ziyang Ma, Southeast University
Valentin Malykh, International IT University
Zhibo Man
Enrique Manjavacas, Sigma Cognition
Marion Marco, Technische Universität München
Guillermo Marco
Evgeny Matusov, AppTek
John McCrae, National University of Ireland Galway
Stephen Meisenbacher
Maggie Mi
Yisong Miao, National University of Singapore
Filip Miletić, University of Stuttgart
Ishani Mondal
Anjishnu Mukherjee, George Mason University
Soichiro Murakami, CyberAgent, Inc.
Arianna Muti
Inderjeet Nair, University of Michigan - Ann Arbor
Pakawat Nakwijit, Queen Mary, University of London
Abhijnan Nath
Vera Neplenbroek
Mariana Neves, German Federal Institute for Risk Assessment
Huy Nghiem
Long Nguyen, Ho Chi Minh city University of Science, Vietnam National University
Lilian Ngweta, IBM Research
Iftitahu Ni'mah, BRIN Indonesia and Eindhoven University of Technology
Irina Nikishina
Dmitry Nikolaev, University of Manchester
Shu Okabe, Technische Universität München
Juri Opitz, University of Zurich
Matthias Orlikowski, Universität Bielefeld
Naoki Otani, Megagon Labs
Kun Ouyang
Valeria Paiva, University of Birmingham
Alexander Panchenko, S-NLP group
Letitia Parcalabescu, Aleph Alpha Research
Tanmay Parekh
Branislav Pecher, Kempelen Institute of Intelligent Technologies
Ted Pedersen, University of Minnesota - Duluth
Davide Picca
Akshara Prabhakar, Salesforce AI Research
Aniket Pramanick, NEC and Technische Universität Darmstadt
Sukannya Purkayastha, NEC
Rifki Putri, Universitas Gadjah Mada
Kun Qian
Shuofei Qiao
Zihan Qiu, Alibaba Group
Lisa Raithel, Technische Universität Berlin
Jishnu Ray Chowdhury, Bloomberg

Frederick Riemenschneider, Ruprecht-Karls-Universität Heidelberg
Anthony Rios, University of Texas at San Antonio
Mathieu Roche, Centre de coopération internationale en recherche agronomique pour le développement
Federico Ruggeri, University of Bologna
Irene Russo, Consiglio Nazionale delle Ricerche
Elisei Rykov, CompSem group
Susanna Rücker, Humboldt-Universität zu Berlin
Yusuke Sakai, Nara Institute of Science and Technology
Xabier Saralegi
Yves Scherrer, University of Oslo
Wesley Scivetti
Anudeex Shetty, University of Melbourne
Ning Shi, University of Alberta
Jisu Shin, Korea Advanced Institute of Science & Technology
KaShun Shum
Yejin Son, Yonsei University
Yueqi Song, Carnegie Mellon University
Richard Sproat, Sakana AI
Shane Storks, University of Michigan - Ann Arbor
Qiushi Sun, University of Hong Kong
Si Sun, Tsinghua University, Tsinghua University
Marek Suppa, Comenius University in Bratislava
Liling Tan
Xingwei Tan, University of Sheffield
Jörg Tiedemann, University of Helsinki
Evgeniia Tokarchuk, University of Amsterdam
MeiHan Tong
Elena Tutubalina
Can Udomcharoenchaikit, Vidyasirimedhi Institute of Science and Technology
Takehito Utsuro, University of Tsukuba
Ashwini Vaidya
Sowmya Vajjala, National Research Council Canada
Eva Vecchi, University of Stuttgart
David Vilar, Google
Emilio Villa-Cueva
Mengru Wang
Qingyun Wang, College of William and Mary
Qiqi Wang, Nankai University
Shuting Wang, Renmin University of China
Taro Watanabe, Nara Institute of Science and Technology
Bonnie Webber, University of Edinburgh
Shengqiong Wu
Yunze Xiao
Rui Xing, Mohamed bin Zayed University of Artificial Intelligence and University of Melbourne
Shanshan Xu, University of Copenhagen
Yige Xu, Nanyang Technological University
Wei Xue
Changbing Yang, University of British Columbia
Zhiyu Yang
Peiran Yao, Amazon

Zhangyue Yin
Zheng Xin Yong, Brown University
Olga Zamaraeva, Universidad de La Coruña
Jingjie Zeng
Haiqi Zhang, University of Texas at Arlington
Ruo Chen Zhang, Brown University
Zheyuan Zhang
Xiutian Zhao, Johns Hopkins University
Yang Zhong, University of Pittsburgh
Naitian Zhou, University of California, Berkeley
Mingyang Zhou, Shenzhen University
Junnan Zhu, Chinese Academy of Sciences
Elena Zotova, Fundación Vicomtech
Pierre Zweigenbaum, LISN, CNRS, Université Paris-Saclay
Rik van Noord, University of Groningen
Elena Álvarez-Mellado, Universidad Nacional de Educación a Distancia
Michal Štefánik, National Institute of Informatics

Table of Contents

<i>EACL 2026 Student Research Workshop: Mentorship Program Report</i> Selene Báez Santamaría, Sai Ashish Somayajula and Atsuki Yamaguchi	1
<i>Mask What Matters: Mitigating Object Hallucinations in Multimodal Large Language Models with Object-Aligned Visual Contrastive Decoding</i> Boqi Chen, Xudong Liu and Jianing Qiu	9
<i>Domain Adaptation of Image Encoder for Multimodal Manga Translation</i> Kota Manabe, Tomoyuki Kajiwara, Takashi Ninomiya, Isao Goto, Shonosuke Ishiwatari and Hiroshi Noji	17
<i>Do Multi-Agents Solve Better Than Single? Evaluating Agentic Frameworks for Diagram-Grounded Geometry Problem Solving and Reasoning</i> Mahbub E Sobhani, Md. Faiyaz Abdullah Sayeedi, Mohammad Nehad Alam, Prama Hossain Progga and Swakkhar Shatabda	27
<i>Luth: Efficient French Specialization for Small Language Models and Cross-Lingual Transfer</i> Maxence Lasbordes and Sinoué Gad	48
<i>Machine Translation for Low-Resource Languages through Monolingual Data and LLM: A Case Study of English-to-Basque</i> Nam Luu, Aitor Soroa, German Rigau and Ondřej Bojar	60
<i>Rethinking the Evaluation of Alignment Methods: Insights into Diversity, Generalisation, and Safety</i> Denis Janiak, Julia Moska, Dawid Motyka, Karolina Seweryn, Paweł Walkowiak, Bartosz Żuk and Arkadiusz Janz	92
<i>Modality Matching Matters: Calibrating Language Distances for Cross-Lingual Transfer in URIEL+</i> York Hay Ng, Aditya Khan, Xiang Lu, Matteo Salloum, Michael Zhou, Phuong Hanh Hoang, A. Seza Dođruöz and En-Shiun Annie Lee	110
<i>Is He Extroverted? Identifying Missing Relevant Personas for Faithful User Simulation</i> Weiwen SU, Yuhan Zhou, Zihan Wang, Naoki Yoshinaga and Masashi Toyoda	131
<i>Quality-Aware Adversarial Ensemble for Singer Identification in 1960s Tamil Film Music</i> Sathiyakugan Balakrishnan and Uthayasanker Thayasivam	150
<i>Thesis Proposal: Efficient KV Cache Reuse for Multi-Document Retrieval-Augmented Generation</i> Zhipeng Zhang and Dmitry Ilvovsky	160
<i>Thesis proposal: COGNILENS: Analyzing Cognitive Decline in Language Models for Alzheimer’s Monitoring</i> Jonathan Guerne	170
<i>Beyond One-Step Distillation: Bridging the Capacity Gap in Small Language Models via Multi-Step Knowledge Transfer</i> Gaeun Yim, Nayoung Ko and Manasa Bharadwaj	182
<i>Thesis proposal: Are We Losing Textual Diversity to Natural Language Processing?</i> Josef Jon and Ondřej Bojar	188
<i>Constructing a Dataset for Hallucination Detection in Japanese Summarization with Fine-grained Faithfulness Labels</i> Hikari Tanaka, Atsushi Keyaki and Mamoru Komachi	207

<i>Comparing Text Compression Capabilities of Large Language Models with Traditional Compression Algorithms</i>	
Mehran Haddadi and William John Teahan	219
<i>Comprehensive Comparison of RAG Methods Across Multi-Domain Conversational QA</i>	
Klejda Alushi, Jan Strich, Chris Biemann and Martin Semmann	233
<i>LEMUR: Robust Fine-Tuning for Multilingual Embedding Models for Retrieval</i>	
Narges Baba Ahmadi, Jan Strich, Martin Semmann and Chris Biemann	248
<i>Trainable, Multiword-aware Linguistic Tokenization Using Modern Neural Networks</i>	
Clara Boesenberg and Kilian Evang	266
<i>Different Time, Different Language: Revisiting the Bias Against Non-Native Speakers in GPT Detectors</i>	
Adnan Al Ali, Jindřich Helcl and Jindřich Libovický	277
<i>Call, Reward, Repeat: Advancing Dialog State Tracking with GRPO and Function Calling</i>	
Timur Ionov, Anna Marshalova and Valentin Malykh	292
<i>Generalising LLM Routing using Past Performance Retrieval: A Few-Shot Router is Sufficient</i>	
Clovis Varangot-Reille, Christophe Bouvard and Antoine Gourru	304
<i>CAPID: Context-Aware PII Detection for Question-Answering Systems</i>	
Mariia Ponomarenko, Sepideh Abedini, Masoumeh Shafieinejad, D. B. Emerson, Shubhankar Mohapatra and Xi He	320
<i>Exploring the Semantic Space of Second Language Learners</i>	
Trisha Godara, Rui He, Wolfram Hinzen and Yan Cong	332
<i>Kahaani: A Multimodal Co-Creative Storytelling System</i>	
Samee Arif, Muhammad Saad Haroon, Aamina Jamal Khan, Taimoor Arif, Agha Ali Raza and Awais Athar	347
<i>A Benchmark and Evaluation of Automated Language of Study Extraction from Computational Linguistics Publications</i>	
Henry Gagnier and Ashwin Kirubakaran	366
<i>Who Plays Which Role? Protagonist Detection and Classification in Moral Discourse</i>	
Mirko Sommer and Maria Becker	375
<i>Thesis Proposal: Multimodal Benchmark for Music Understanding in Large Language Models</i>	
Tomáš Sourada	393
<i>Communication as a Complex System: Modeling the Feedback Dynamics of Trust and Credibility</i>	
Swaptik Chowdhury, Samuel D. Allen and Jung Hee Hyun	406
<i>The Clinical Fingerprint: Comparing the Rhetorical Integrity and Epistemic Safety of Human Physicians and Large Language Models</i>	
Bayram Ayadi	416
<i>Acceleration of Backpropagation in Linear Layers of Transformer Models Based on Gradient Structure</i>	
Dmitrii Topchii, Alexander Panchenko and Viktoriia A. Chekalina	426
<i>Chronocept: Instilling a Sense of Time in Machines</i>	
Krish Goel, Sanskar Pandey, Mahadevan KS, Harsh Kumar and Vishesh Khadaria	437
<i>When Prompt Optimization Becomes Jailbreaking: Adaptive Red-Teaming of Large Language Models</i>	
Zafir Shamsi, Nikhil Chekuru, Zachary Guzman and Shivank Garg	457

<i>GraphRAG-Rad: Concept-Aware Radiology Report Generation via Latent Visual-Semantic Retrieval</i> Faezeh Safari, Hang Dong, Zeyu FU and Aline Villavicencio	464
<i>Token Pruning for Improving Graph-Generating State Space Model Performance</i> Monish Beegamudre, Jack Zheng and Margaret Capetz	476
<i>Scale Is All You Need: Analyzing Modality Interaction and Speaker Intent Without Fine-Tuning</i> Animesh Gurjar and Nikhil Krishnaswamy	483
<i>Plasticity vs. Rigidity: The Impact of Low-Rank Adapters on Reasoning on a Micro-Budget</i> Zohaib Khan, Omer Tafveez and Zoha Hayat Bhatti	493
<i>In-Image Machine Translation. A Preliminary Modular Approach</i> Sergio Gomez Gonzalez, Miguel Domingo and Francisco Casacuberta	502
<i>Text-to-Text Automatic Story Generation: A Survey</i> Yuan Ma, Hanna Suominen, Patrik Haslum and Richard Susilo	514
<i>Probabilistic Bilingual Subword Segmentation with Latent Subword Alignment</i> Shoto Nishida, Daiki Matsui, Takashi Ninomiya, Isao Goto and Akihiro Tamura	528
<i>Thesis Proposal: Development of End-to-End Speech Translation Models for Indian Languages</i> Jamaluddin	535
<i>Towards Singable Lyrics Translation Using Large Language Models</i> Liu Hanze, Yusuke Sakai and Taro Watanabe	544
<i>Evaluating the Impact of SAE-based Language Steering on LLM Performance</i> Sebastian Zwirner, Wentao Hu, Koshiro Aoki and Daisuke Kawahara	555
<i>Annotation-Efficient Vision-Language Model Adaptation to the Polish Language Using the LLaVA Framework</i> Grzegorz Statkiewicz, Alicja Dobrzeniecka, Karolina Seweryn, Aleksandra Krasnodebska, Karolina Piosek, Katarzyna Bogusz, Sebastian Cygert and Wojciech Kusa	569
<i>A Computational Forensic Linguistic Analysis of Narrative and Question-Answer Structures in Italian Police Interrogation Transcripts</i> Romane Werner, Thomas François and Sonja Bitzer	590
<i>Thesis Proposal: A Multi-Agent System for Ontology-Based Perspective-Aware Knowledge Extraction</i> Luiz Do Valle Miranda and Grzegorz J. Nalepa	604
<i>Fake News Detection Strategies under Dataset Bias: Using Large-scale Coarse-grained Labels</i> Yuki Kishi, Yuji Arima and Hitoshi Iyatomi	612
<i>DRAGOn: Designing RAG On Periodically Updated Corpus</i> Fedor Chernogorskii, Sergei Averkiev, Liliya Kudrалеeva, Zaven Martirosian, Maria Tikhonova, Valentin Malykh and Alena Fenogenova	622
<i>Efficient Low-Resource Language Models Using Tokenizer Transfer</i> Gustaf Gren and Murathan Kurfali	639
<i>Learning Nested Named Entity Recognition from Flat Annotations</i> Igor Rozhkov and Natalia V Loukachevitch	649
<i>Analysing LLM Persona Generation and Fairness Interpretation in Polarised Geopolitical Contexts</i> Maida Aizaz and Quang Minh Nguyen	664

<i>Beyond Bias Scores: Unmasking Vacuous Neutrality in Small Language Models</i> Sumanth Manduru and Carlotta Domeniconi	685
<i>From Detection to Explanation: Modeling Fine-Grained Emotional Social Influence Techniques with LLMs and Human Preferences</i> Maciej Markiewicz, Wiktoria Mieleszczenko-Kowszewicz, Beata Bajcar, Tomasz Adamczyk, Aleksander Szczęsny, Jolanta Babiak and Przemysław Kazienko	715
<i>How Do Lexical Senses Correspond Between Spoken German and German Sign Language?</i> Melis Çelikkol and Wei Zhao	735
<i>Evaluating Cost-Efficiency of LLMs in a RAG Setup on Polish Wikipedia: Quality vs. Energy Consumption</i> Patrycja Smits and Tomasz Walkowiak	747
<i>Thesis Proposal: Measuring Prejudice at Scale</i> Zoran Fijavž, Senja Pollak and Veronika Bajt	760
<i>Energy Matching based Preference Learning for Diffusion Language Models</i> Shiv Shankar	776
<i>Thesis Proposal: Stability-Aware, Evidence-Grounded Knowledge Graph for Substance Use Disorders and Social Determinants of Health</i> Gautham Vijay Kumar	787
<i>Detecting Overflow in Compressed Token Representations for Retrieval-Augmented Generation</i> Julia Belikova, Danila Rozhevskii, Dennis Svirin, Konstantin Polev and Alexander Panchenko	797
<i>Irnrx: A library for Linear RNNs</i> Karan Bania, Soham Kalburgi, Manit Tanwar, Dhruthi, Aditya Nagarsekar, Harshvardhan Mestha, Naman Chibber, Raj Deshmukh, Anish Sathyanarayanan, Aarush Rathore and Pratham Chheda . . .	811
<i>Automatic Generation of a Compositional QA Benchmark for Geospatial Reasoning under Spatial and Entity Constraints</i> Tetsuhisa Suizu, Shohei Higashiyama, Hiroyuki Shindo, Hiroki Ouchi and Sakriani Sakti . . .	818
<i>Thesis Proposal: Comparing Human and Model Perception of Writing Style under Controlled Perturbations</i> Ewelina Paulina Księżniak	831
<i>Bring the Apple, Not the Sofa: Impact of Irrelevant Context in Embodied AI Commands on VLA Models</i> Andrey Moskalenko, Daria Pugacheva, Denis Shepelev, Andrey Kuznetsov, Vlad Shakhuro and Elena Tutubalina	840
<i>An Evaluation of Classifiers for Mapping Generative LLM Responses to Answer Options of Multiple-choice Questionnaires</i> Alisea Stroligo, Anna Shamray, Julian Schelb and Andreas Spitz	861
<i>Pioneering Bot Detection on Polish Reddit at the Comment Level</i> Karmela Matyjaszek	881
<i>What the Router Sees Matters: Funnel Pooling for Fast, Content Driven Expert Routing</i> Josef Pichlmeier, Sebastian Nicolas Mueller, Jakob Sturm, Josef Dräxl and Andre Luckow . .	895
<i>TimeRes: A Turkish Benchmark For Evaluating Temporal Understanding of Large Language Models</i> Habib Yağız Demir, Ümit Atlamaz and Susan Üsküdarlı	910

<i>Hospitality-VQA: Decision-Oriented Informativeness Evaluation for Vision–Language Models</i> Jeongwoo Lee, Baek Duhyeong, Eungyeol Han, Soyeon Shin, Gukin Han, Seungduk Kim, Jae- hyun Jeon and Taewoo Jeong	921
<i>Colorism in Multimodal AI: An Empirical Exploration of Socioeconomic Linguistic Bias in Text-to- Image Generation</i> Raj Gaurav Maurya, Vaibhav Shukla and Sreedath Panat	937
<i>Active Learning for Corpus Refinement: Cost-Effective Preprocessing to Improve Validity of Applied Quantitative Text Analysis</i> Jakob Steglich and Stephan Poppe	952
<i>From Sentences to Proof Trees: Leveraging Language Models for Structured Reasoning</i> Aayushee Gupta	967

Program

Thursday, March 26, 2026

10:30 - 09:00 *Opening Remarks, Best Paper Award Presentation, and Spotlight Talks*

10:30 - 11:00 *Break*

11:00 - 12:30 *Poster Session (6)*

12:30 - 14:30 *Break*

14:30 - 16:00 *Poster Session (8)*