

EACL 2026

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

Proceedings of System Demonstrations

March 24-29, 2026

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

Platinum



Megagon Labs

Appen

Bronze

elra

 translated.



©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 979-8-89176-382-1

Introduction

Welcome to the proceedings of the System Demonstration track of the 19th Conference of the European Chapter of the Association for Computational Linguistics (EACL 2026), held in Rabat, Morocco, from March 24 to 29, 2026.

For the EACL 2026 System Demonstration track, we received 102 submissions, almost doubling the number of submissions compared to the previous edition. Of these, 44 were selected for inclusion in the program (acceptance rate of 43.1%). The selection process was highly competitive, and many strong submissions could unfortunately not be accommodated.

We would like to thank the members of the program committee and area chairs for their careful and timely reviews. We are also grateful to the organizing and publication teams for their support. Finally, we sincerely thank all authors who submitted their work to the Demonstrations track and contributed to the scientific quality of the program.

The accepted demonstration papers reflect the breadth and vitality of current research in Natural Language Processing, covering evaluation frameworks, interpretability tools, dialogue systems, knowledge-based applications, and responsible AI platforms. The demonstrations will be presented in person during the conference, including dedicated poster and demo sessions.

We appreciate the efforts made by all authors to showcase their work and contribute to the success of EACL 2026.

Danilo Croce, Jochen L. Leidner, Nafise Sadat Moosavi
EACL 2026 System Demonstration Chairs

Program Committee

Chairs

Danilo Croce, University of Roma Tor Vergata
Jochen L. Leidner, KnowledgeSpaces UG (haftungsbeschränkt), Coburg University of Applied Sciences and University of Sheffield
Nafise Sadat Moosavi, University of Sheffield

Area Chairs

Hiba Arnaout, Technische Universität Darmstadt
Long Bai, Institute of Computing Technology, Chinese Academy of Sciences
Simone Balloccu, Technische Universität Darmstadt
Valerio Basile, University of Turin
Giuseppe Castellucci, Amazon
Craig Erickson, Amazon
Junxian He, Hong Kong University of Science and Technology
Christian Heumann, Ludwig-Maximilians-Universität München
Shaoxiong Ji, University of Turku and ELLIS Institute Finland
Khalid Al Khatib, University of Groningen
Christopher Klamm, University of Cologne and University of Mannheim
Anne Lauscher, Universität Hamburg
Kelly Marchisio, Cohere and Cohere
John Philip McCrae, National University of Ireland Galway
Qingyun Wang, College of William and Mary
Hua Wei, Arizona State University

Reviewers

Angus Addlesee, Amazon
Suman Adhya, Indian Association for the Cultivation of Science
Yamen Ajjour, Universität Hannover
Seyed Alireza Mousavian Anaraki, University of Roma Tor Vergata
Miriam Anschütz, Technische Universität München
Mario Ezra Aragon, Universidad de Santiago de Compostela
Matthias Aßenmacher, Ludwig-Maximilians-Universität München
Joanne Boisson, Cardiff University
Federico Borazio, University of Roma Tor Vergata
Georgeta Bordea, La Rochelle Université
Marco Braga, University of Milan - Bicocca
Marco Antonio Sobrevilla Cabezudo, Pontificia Universidad Católica del Perú
Aunabil Chakma, University of Arizona
Raghuveer Chanda, Google
Nischal Reddy Chandra, Adobe Systems
Tiejin Chen, Arizona State University
Daiwei Chen, University of Wisconsin - Madison
Chung-Chi Chen, AIST, National Institute of Advanced Industrial Science and Technology
Kehai Chen, Harbin Institute of Technology (Shenzhen)

Sanjay Das, University of Texas at Dallas
Ona de Gibert, University of Helsinki
Giacomo De Luca, University of Roma Tor Vergata
Jacob Devasier, University of Texas at Arlington
Luigi Di Caro, University of Turin
Kshitij P Fadnis, IBM Research
Enfa Fane, University of Arizona
Mariia Fedorova, University of Oslo
Nils Feldhus, Technische Universität Berlin
Tim Fischer, University of Hamburg
Esteban Garces Arias, Ludwig-Maximilians-Universität München
Javier García Gilabert, Barcelona Supercomputing Center
Tirthankar Ghosal, Oak Ridge National Laboratory
Voula Giouli, Aristotle University of Thessaloniki and ILSP - AthenaResearch Center
George Gkotsis, Upwork Inc
Xuehang Guo, College of William and Mary
Claudiu Daniel Hromei, University of Roma Tor Vergata
Dayu Hu, Northeastern University
Timour Igamberdiev, Universität Vienna
Oghenevovwe Ikumariogbe, University of Arizona
Hasan Iqbal, Mohamed bin Zayed University of Artificial Intelligence
Vidhyakshaya Kannan, Georgia Institute of Technology
Yoshihide Kato, Nagoya University
Jeonghwan Kim, University of Illinois at Urbana-Champaign
Andrei Kucharavy, University of Applied Sciences Western Switzerland, Sierre (HES-SO Valais)
Leo Leppänen, University of Helsinki
Weijiang Li, University of Notre Dame
Zixuan Li, Institute of Computing Technology, Chinese Academy of Sciences
Ruo Chen Li, University of Texas at Dallas
Hanming Li, Tsinghua University
Xiaochang Li, William and Mary
Shuhang Lin, , Rutgers University
Yueqian Lin, Duke University
Ke Lin, College of William and Mary
Zhexiong Liu, University of Pittsburgh
Junteng Liu, The Hong Kong University of Science and Technology
Wenxuan Liu, Chinese Academy of Sciences
Tianrui Liu, Apple
Xinyu Lu, Chinese Academy of Sciences
Maria Mahbub, Oak Ridge National Laboratory
Puneet Mathur, Adobe Systems
Tim Menzner, Hochschule Coburg
Yasmin Moslem, Trinity College Dublin
Youyang Ng, Kioxia Corporation
Tuan-Phong Nguyen, VNU University of Engineering and Technology
Huy V. Nguyen, Amazon
Leyi Pan, Tsinghua University
Jungyeul Park, Korea Advanced Institute of Science & Technology
Max Ploner, Humboldt Universität Berlin
Cheng Qian, University of Illinois at Urbana-Champaign
Yide Ran, Stevens Institute of Technology

Prajvi Saxena, German Research Center for AI
Carsten Schnober, Netherlands eScience Center
Tim Schopf, Technische Universität Dresden
Seongbum Seo, Sejong University
Alay Dilipbhai Shah, Together AI
KaShun Shum, Hong Kong University of Science and Technology
Gosuddin Kamaruddin Siddiqi, Microsoft
Jyotika Singh, Oracle
Bhavyaajeet Singh, Technische Universität Darmstadt
Jiahe Song, Shanghai Jiaotong University and Shanghai Artificial Intelligence Laboratory
Daniil Sorokin, Amazon Development Center Germany
Dachun Sun, Worcester Polytechnic Institute
Qiang Sun, University of Western Australia
Marek Suppa, Comenius University in Bratislava
Sotaro Takeshita, University of Technology Nuremberg and Universität Mannheim
Yudong Tao, Facebook
Jonathan Tonglet, Technische Universität Darmstadt
Vatsal Venkatkrishna, Institute for Computer Science, Artificial Intelligence and Technology
Sami Virpioja, University of Helsinki
Doan Nam Long Vu, Technische Universität Darmstadt
Chuan-Ju Wang, Academia Sinica
Zibu Wei, University of California
Liang Xie, Guangdong University of Technology
Shuo Xing, Google and Texas A&M University - College Station
Xin Xu, The Hong Kong University of Science and Technology
Bingbing Xu, Renmin University of China
Hao Yu, National University of Defense Technology
Pengfei Yu, Amazon
Yige Yuan, University of Washington
Linfan Zhang, University of California
Chongsheng Zhang, Ludwig-Maximilians-Universität München and Henan University
Tong Zhou, Chinese Academy of Science
Zichen Zhu, Shanghai Jiao Tong University
Vilém Zouhar, ETHZ - ETH Zurich
Dennis Zyska, Technische Universität Darmstadt

Table of Contents

<i>Stakeholder Suite: A Unified AI Framework for Mapping Actors, Topics and Arguments in Public Debates</i>	
Mohamed Chenene, Jeanne Rouhier, Jean Daniélou, Mihir Sarkar and Elena Cabrio	1
<i>DeepPavlov Strikes Back: A Toolkit for Improving LLM Reliability and Trustworthiness</i>	
Evgenii Nikolaev, Timur Ionov, Anna Korzanova, Vasily Konovalov and Maksim Savkin	21
<i>PropGenie: A Multi-Agent Conversational Framework for Real Estate Assistance</i>	
Chang Shen, Shaozu Yuan, Kuizong Wu, Long Xu and Meng Chen	33
<i>Pro-QuEST: A Prompt-chain based Quiz Engine for testing Specialized Technical Product Knowledge</i>	
Sujatha Das Gollapalli, Mouad Hakam, Mingzhe Du, See-Kiong Ng and Mohammed Hamzeh	46
<i>elfen: A Python Package for Efficient Linguistic Feature Extraction for Natural Language Datasets</i>	
Maximilian Maurer	61
<i>DELTA: A Toolkit for Measuring Linguistic Diversity in Dependency-Parsed Corpora</i>	
Louis Estève and Kaja Dobrovoljc	75
<i>CLARIESG: An End-to-End System for ESG Analysis over Complex Tables in Corporate Reports</i>	
Marta Santacroce, Michele Luca Contalbo, Sara Pederzoli, Riccardo Benassi, Venturelli Valeria, Matteo Paganelli and Francesco Guerra	86
<i>Fact Finder - Enhancing Domain Expertise of Large Language Models by Incorporating Knowledge Graphs</i>	
Daniel Steinigen, Roman Teucher, Timm Heine Ruland, Max Rudat, Nicolas Flores-Herr, Peter Fischer, Nikola Milosevic, Christopher Schymura and Angelo Ziletti	101
<i>Simplifying Outcomes of Language Model Component Analyses with ELIA</i>	
Aaron Louis Eidt and Nils Feldhus	111
<i>IntelliCode: A Multi-Agent LLM Tutoring System with Centralized Learner Modeling</i>	
Jones David and Shreya Ghosh	129
<i>FiMMIA: scaling semantic perturbation-based membership inference across modalities</i>	
Anton Emelyanov, Sergei Kudriashov and Alena Fenogenova	139
<i>A Browser-based Open Source Assistant for Multimodal Content Verification</i>	
Rosanna Milner, Michael Foster, Olesya Razuvayevskaya, Valentin Porcellini, Denis Teyssou, Ian Roberts and Kalina Bontcheva	154
<i>Infherno: End-to-end Agent-based FHIR Resource Synthesis from Free-form Clinical Notes</i>	
Johann Frei, Nils Feldhus, Lisa Raithel, Roland Roller, Alexander Meyer and Frank Kramer .	163
<i>BOOM: Beyond Only One Modality KIT's Multimodal Multilingual Lecture Companion</i>	
Sai Koneru, Fabian Retkowski, Christian Huber, Lukas Hilgert, Seymanur Akti, Enes Yavuz Ugan, Alexander Waibel and Jan Niehues	175
<i>PEFT-Factory: Unified Parameter-Efficient Fine-Tuning of Autoregressive Large Language Models</i>	
Robert Belanec, Ivan Srba and Maria Bielikova	188
<i>Similar, but why? A Toolkit for Explaining Text Similarity</i>	
Juri Opitz, Andrianos Michail, Lucas Moeller, Sebastian Padó and Simon Clematide	203

<i>AlignFix: A Tool for Parallel Corpora Augmentation and Refinement</i> Samuel Frontull and Simon Haller-Seeber	215
<i>PromptLab: A Collaborative Platform for Prompt Engineering and Dataset Curation</i> Maged S. Al-shaibani, Zaid Alyafeai, Dania Refai, Nawaf Alomari, Ahmed Ashraf, Mais Alheraki, Mustafa Alturki, Hamzah Luqman and Irfan Ahmad	225
<i>LLM BiasScope: A Real-Time Bias Analysis Platform for Comparative LLM Evaluation</i> Himel Ghosh and Nick Elias Werner	261
<i>InkSight: Towards AI-Aided Historical Manuscript Analysis</i> Andrey Sakhovskiy, Ivan Ulitin, Emilia Bojarskaja, Vladimir Kokh, Ruslan Murtazin, Maxim Novopoltsev and Semen Budenny	271
<i>promptlotion: A Unified, Modular Framework for Prompt Optimization</i> Tom Zehle, Timo Heiß, Moritz Schlager, Matthias Aßenmacher and Matthias Feurer	282
<i>T-pro 2.0: An Efficient Russian Hybrid-Reasoning Model and Playground</i> Dmitrii Stoianov, Danil Taranets, Olga Tsymboi, Ramil Latypov, Almaz Dautov, Vladislav Kru- glikov, Surkov Nikita, German Abramov, Pavel Gein, Dmitry Abulkhanov, Mikhail Gashkov, Viktor Zelenkovskiy, Artem Batalov, Aleksandr Medvedev and Anatolii Potapov	297
<i>SDialog: A Python Toolkit for End-to-End Agent Building, User Simulation, Dialog Generation, and Evaluation</i> Sergio Burdisso, Séverin Baroudi, Yanis Labrak, David Grünert, Pawel Cyrt, Yiyang Chen, Srikanth Madikeri, Esaú Villatoro-tello, Ricard Marxer and Petr Motlicek	320
<i>Agentic AI for Human Resources: LLM-Driven Candidate Assessment</i> Kamer Ali Yuksel, Abdul Basit Anees, Ashraf Hatim Elneima, Sanjika Hewavitharana, Mohamed Al-Badrashiny and Hassan Sawaf	341
<i>Trove: A Flexible Toolkit for Dense Retrieval</i> Reza Esfandiarpour, Max Zuo and Stephen Bach	349
<i>ClinicalTrialsHub: Bridging Registries and Literature for Comprehensive Clinical Trial Access</i> Jiwoo Park, Ruoqi Liu, Avani Jagdale, Andrew Srisuwananukorn, Jing Zhao, Ping Zhang and Sachin Kumar	359
<i>SciTrue: Evidence-Grounded Claim Verification in Science</i> Neset Tan, Minghao LI and Mark Gahegan	397
<i>PUCP-Matrix: An Open-source and Comprehensive Toolkit for Linguistic Analysis of Spanish Texts</i> Javier Alonso Villegas Luis and Marco Antonio Sobrevilla Cabezudo	407
<i>Integrity Shield A System for Ethical AI Use & Authorship Transparency in Assessments</i> Ashish Raj Shekhar, Shiven Agarwal, Priyanuj Bordoloi, Yash Shah, Tejas Anvekar and Vivek Gupta	417
<i>Using a Human-AI Teaming Approach to Create and Curate Scientific Datasets with the SciLire System</i> Necva Bölücü, Jessica Irons, Changhyun Lee, Brian Jin, Maciej Rybinski, Huichen Yang, Andreas Duenser and Stephen Wan	428
<i>xLM: A Python Package for Non-Autoregressive Language Models</i> Dhruvesh Patel, Durga Prasad Maram, Sai Sreenivas Chintha, Benjamin Rozonoyer and Andrew McCallum	445

<i>AITutor-EvalKit: Exploring the Capabilities of AI Tutors</i>	
Numaan Naeem, Kaushal Kumar Maurya, Kseniia Petukhova and Ekaterina Kochmar	457
<i>EvalSense: A Framework for Domain-Specific LLM (Meta-)Evaluation</i>	
Adam Dejl and Jonathan Pearson	480
<i>AI for Climate Finance: Agentic Retrieval and Multi-Step Reasoning for Early Warning System Investments</i>	
Saeid Vaghefi, Aymane Hachcham, Veronica Grasso, Nakiete Msemu, Chiara Colesanti Senni and Markus Leippold	492
<i>RAGVUE: A Diagnostic View for Explainable and Automated Evaluation of Retrieval-Augmented Generation</i>	
Keerthana Murugaraj, Salima Lamsiyah and Martin Theobald	512
<i>SmartMatch: Real-Time Semantic Retrieval for Translation Memory Systems</i>	
Ernesto Luis Estevanell Valladares, Salima Lamsiyah, Alicia Picazo-Izquierdo, Tharindu Ranasinghe, Ruslan Mitkov and Rafael Munoz	527
<i>QSTN: A Modular Framework for Robust Questionnaire Inference with Large Language Models</i>	
Maximilian Kreutner, Jens Rupperecht, Georg Ahnert, Ahmed Salem and Markus Strohmaier	537
<i>A Virtual Assistant for Architectural Design in a VR Environment</i>	
Ander Salaberria, Oier Ijurco, Markel Ferro, Jiayuan Wang, Iñigo Vilá Muñoz, Roberto de Ioris, Jeremy Barnes and Oier Lopez De Lacalle	550
<i>ARGSBASE: A Multi-Agent Interface for Structured Human–AI Deliberation</i>	
Frieso Turkstra, Sara Nabhani and Khalid Al Khatib	563
<i>Simultaneous Speech-to-Text Translation Web Application for Estonian</i>	
Bohdan Podziubanchuk, Aivo Olev, Jiaming Kong and Tanel Alumäe	575
<i>The AI Committee: A Multi-Agent Framework for Automated Validation and Remediation of Web-Sourced Data</i>	
Sunith Vallabhaneni, Thomas Berkane and Maimuna S. Majumder	583
<i>entity-linkings: A Unified Library for Entity Linking</i>	
Yuya Sawada, Tsuyoshi Fujita, Yusuke Sakai and Taro Watanabe	591
<i>ESG-KG: A Multi-modal Knowledge Graph System for Automated Compliance Assessment</i>	
Li-Yang Chang, Chih-Ming Chen, Hen-Hsen Huang, Ming-Feng Tsai, An-Zi Yen and Chuan-Ju Wang	602
<i>BanSuite: A Unified Toolkit and Software Platform for Low-Resource NLP in Bangla</i>	
Md. Abu Sayed, Faisal Ahamed Khan, Jannatul Ferdous Tuli, Nabeel Mohammed, Mohammad Ruhul Amin and Mohammad Mamun Or Rashid	609