

EACL 2026

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

Volume 5: Industry Track

March 25-27, 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 979-8-89176-384-5

Introduction

We are pleased to welcome you to the EACL 2026 Industry Track, held on 25–27 March 2026, during the main conference days of the 19th Conference of the European Chapter of the Association for Computational Linguistics. The track attracted 170 submissions, and a total of 167 reviewers contributed to the review process. After double-blind peer review, 71 papers were selected for presentation at the EACL 2026 Industry Track. Of these, 57 papers will be presented as oral talks (26 virtually) and 14 as posters.

Thematically, large language models are central to many submissions, with a strong emphasis on grounded, production-facing systems rather than standalone generation. Prominent themes include retrieval and grounding, with advances in retrieval-augmented generation, re-ranking, robustness to noisy or redundant retrieval context, and compliance-aware search. Another major thread involves structured and agentic workflows for complex, multi-step tasks, alongside rigorous evaluation via new benchmarks and error-analysis frameworks covering reasoning, information extraction, and safety. Many papers also focus on multimodal understanding, document and table understanding, and domain adaptation under real-world constraints such as latency, privacy, and multilingual settings. These directions span applications in healthcare, finance, law and enterprise, e-commerce and search advertising, workforce analytics, and interactive decision support.

We would like to thank the authors of all Industry Track submissions, as well as the reviewers, for their hard work and dedication under tight deadlines. We also thank the General Chair, the Technical Open Review Chair, the Virtual Infrastructure Chairs, the Website Chair, the Publication Chairs, and all the other EACL 2026 committees. Finally, we would like to thank the ACL team and the Underline team, especially Jennifer Rachford and Damira Mrsic.

Yevgen Matusevych, Nikolaos Aletras, Gülşen Eryiğit
Industry Track Chairs

Program Committee

Industry Track Chairs

Yevgen Matuskevych, University of Groningen, the Netherlands
Gülşen Eryiğit, Istanbul Technical University, Türkiye
Nikolaos Aletras, University of Sheffield, England

Reviewers

Ahmed Abdelali, Hanna Abi Akl, Sallam Abualhaija, Rupam Acharyya, Georgios Alexandridis, Duygu Altinok, Mario Ezra Aragon, Kushagr Arora, Tuğba Pamay Arslan, Ankit Arun

Long Bai, Mithun Balakrishna, Yuwei Bao, Leslie Barrett, Daniel Bauer, Dario Bertero, Arjun Bhalla, Mukul Bhutani, Debmalya Biswas, Nadjat Bouayad-Agha, Quentin Brabant, Thomas Brovelli, Hannah Brown

Fabio Casati, Dumitru-Clementin Cercel, Ankani Chattoraj, Fuxiang Chen, Guanhua Chen, Jian-ning Chen, John Chen, Lin Chen, Yubo Chen, Won Ik Cho, Shamil Chollampatt, Bonaventura Coppola

Deborah A. Dahl, Daniel Dakota, Aswarth Abhilash Dara, Souvik Das, Tirthankar Dasgupta, Steve DeNeefe, Daryna Dementieva, Prajit Dhar, Daniel Dickinson, Rahul Divekar, Bin Dong, Matthew T. Dunn

Aparna Elangovan, Keelan Evanini

Run-Ze Fan, Yihao Fang, Michael Flor, Simona Frenda, Lisheng Fu

Baban Gain, Mozhddeh Gheini, Sucheta Ghosh, Voula Giouli, Vinod Goje, Jiaying Gong, Tong Guo, Ramiro H. Gálvez

Maeda Hanafi, Derrick Higgins, Pengyu Hong, Sho Hoshino, Jiahe Huang, John S Hudzina

Srideepika Jayaraman, Jiyue Jiang, Zhuoxuan Jiang, Shailza Jolly

Anup K. Kalia, Geewook Kim, Tracy Holloway King, Ana Kotarcic, Rajasekar Krishnamurthy, Marek Kubis, Andrei Kucharavy, Sanjeev Kumar, Kemal Kurniawan

Yanis Labrak, Stefan Larson, Md Tahmid Rahman Laskar, Alexandra Lavrentovich, Arun Balajiee Lekshmi Narayanan, Yves Lepage, Chong Li, Keyi Li, Yinghui Li, Yingya Li, Veronica Liesaputra, Gilbert Lim, Dongqi Liu, Lei Liu, Pengfei Liu, Ye Liu, Yonghao Liu, Alejandra Lorenzo, Natalia V Loukachevitch

Liang Ma, Zhixin Ma, Wolfgang Maier, Fred Mailhot, Lorenzo Malandri, Yuji Matsumoto, David D. McDonald, Alexander Mehler, Mahnoosh Mehrabani, Yuanliang Meng, Mohsen Mesgar, Margot Mieskes, Hideya Mino, Hemant Misra, Lori Moon, Matthew Mulholland, Emir Muñoz

Tetsuji Nakagawa, Sungjin Nam, Marcin Namysl, Diane Napolitano, Navid Nobani

Alexander O'Connor, Oleg Okun, Eda Okur, Berke Oral, Naoki Otani

Xueting Pan, Pierre-Henri Paris, Vera Pavlova, Ali Pesaranghader, Jakub Piskorski, Animesh Prasad

Xin Ying Qiu

Shihao Ran, Traian Rebedea, Joy Rimchala, Brian Riordan, Susanna Rücker

Alicia Sagae, Tanay Kumar Saha, Minoru Sasaki, Ulrich Schäfer, Shubhashis Sengupta, Ronald Seoh, Sofia Serrano, Manali Sharma, Tianhao Shen, Qiang Sheng, Ashish Shenoy, Andrew Silva, Mukul Singh, Priyanka Sinha, Hyun-Je Song, Ian Stewart, Kristina Striegnitz, Sebastian Stüker, Marek Suppa, Munira Syed

Xuemei Tang, Khushboo Thaker, Ketan Thakkar, Manabu Torii, Giuliano Tortoreto, Aashka Trivedi, Keith Trnka, Masaaki Tsuchida

Brian Ulicny

Daniel Varab, Manuel Vilares Ferro

Jianzong Wang, Ryan Wang, Penghui Wei, Haryo Akbarianto Wibowo, Tianxing Wu

Kaige Xie, Siheng Xiong

Manoj Yadav, Guanqun Yang, Zhengzhe Yang, Bingyang Ye, Dezhi Ye, Jinyeong Yim, Yuwei Yin, Issei Yoshida, Cheng Yu, Lei Yu

Mahdi Zakizadeh, Lei Zhang, Ningyu Zhang, Yin Zhang, Baohang Zhou, Dong Zhou, Xiliang Zhu, Yuqicheng Zhu, Imed Zitouni

Table of Contents

<i>Iterative Structured Pruning for Large Language Models with Multi-Domain Calibration</i> Guangxin Wu, Hao Zhang, Zhang Zhibin, Jiafeng Guo and Xueqi Cheng	1
<i>SCRIPTMIND: Crime Script Inference and Cognitive Evaluation for LLM-based Social Engineering Scam Detection System</i> Heedou Kim, Changsik Kim, Sanghwa Shin and Jaewoo Kang	11
<i>From Paper to Structured JSON: An Agentic AI Workflow for Compliant BMR Digital Transformation</i> Bhavik Agarwal, Nidhi Bendre and Viktoria Rojkova	39
<i>Compact Multimodal Language Models as Robust OCR Alternatives for Noisy Textual Clinical Reports</i> Nikita Neveditsin, Pawan Lingras, Salil Patil, Swarup Patil and Vijay Kumar Mago	48
<i>PersonaTrace: Synthesizing Realistic Digital Footprints with LLM Agents</i> Minjia Wang, Yunfeng Wang, Xiao Ma, Dexin Lv, Qifan Guo, Lynn Zheng, Benliang Wang, Lei Wang, Jiannan Li, Yongwei Xing, Junzhe Xu and Zheng Sun	60
<i>Evaluating the Pre-Consultation Ability of LLMs using Diagnostic Guidelines</i> Jean Seo, Gibaeg Kim, Kihun Shin, Seungseop Lim, Hyunkyung Lee, Wooseok Han, Jongwon Lee and Eunho Yang	78
<i>SELENE: Selective and Evidence-Weighted LLM Debating for Efficient and Reliable Reasoning</i> Akshay Verma, Swapnil Gupta, Deepak Gupta, Prateek Sircar and Siddharth Pillai	95
<i>SymPyBench: A Dynamic Benchmark for Scientific Reasoning with Executable Python Code</i> Shima Imani, Seungwhan Moon, Adel Ahmadyan, Lu Zhang, Ahmed Kirmani and Babak Damavandi	105
<i>KV Pareto: Systems-Level Optimization of KV Cache and Model Compression for Long Context Inference</i> Sai Gokhale, Devleena Das, Rajeev Patwari, Ashish Sirasao and Elliott Delaye	119
<i>MizanQA: A Benchmark for Multi-Answer Moroccan Legal QA</i> Adil Bahaj and Mounir Ghogho	132
<i>Router-Suggest: Dynamic Routing for Multimodal Auto-Completion in Visually-Grounded Dialogs</i> Sandeep Mishra, Devichand Budagam, Anubhab Mandal, Bishal Santra, Pawan Goyal and Manish Gupta	145
<i>Beyond Unified Models: A Service-Oriented Approach to Low Latency, Context Aware Phonemization for Real Time TTS</i> Mahta Fetrat Qharabagh, Donya Navabi, Zahra Dehghanian, Morteza Abolghasemi and Hamid R. Rabiee	157
<i>Retrieval Enhancements for RAG: Insights from a Deployed Customer Support Chatbot</i> Daniel González Juclà, Mohit Tuteja, Marcos Esteve Casademunt, Keshav Unnikrishnan, Yasir Usmani and Arvind Roshaan	169
<i>Scaling Intent Understanding: A Framework for Classification with Clarification using Lightweight LLMs</i> Subhadip Nandi, Tanishka Agarwal, Anshika Singh and Priyanka Bhatt	181
<i>Beyond IVR: Benchmarking Customer Support LLM Agents for Business-Adherence</i> Sumanth Balaji, Piyush Mishra, Aashraya Sachdeva and Suraj Agrawal	193

<i>HotelQuEST: Balancing Quality and Efficiency in Agentic Search</i>	
Guy Hadad, Shadi Iskander, Sofia Tolmach, Oren Kalinsky, Haggai Roitman and Ran Levy .	209
<i>TASER: Table Agents for Schema-guided Extraction and Recommendation</i>	
Nicole Cho, Kirsty Fielding, William Watson, Sumitra Ganesh and Manuela Veloso	226
<i>TAGQuant: Token-Aware Clustering for Group-Wise Quantization</i>	
Jaeseong Lee, Seung-won Hwang, Aurick Qiao, Zhewei Yao and Yuxiong He	253
<i>Beyond Grid Search: Leveraging Bayesian Optimization for Accelerating RAG Pipeline Optimization</i>	
Anum Afzal, Xueru Zheng and Florian Matthes	263
<i>BornoDrishti: Leveraging Vision Encoders and Domain-Adaptive Learning for Bangla OCR on Diverse Documents</i>	
S M Jishanul Islam, Md Mehedi Hasan, Masbul Haider Ovi, Akm Shahariar Azad Rabby and Fuad Rahman	278
<i>MobileCity: An Efficient Framework for Large-Scale Urban Behavior Simulation</i>	
Xiaotong Ye, Nicolas Bougie, Toshihiko Yamasaki and Narimawa Watanabe	287
<i>Is Micro Domain-Adaptive Pre-Training Effective for Real-World Operations? Multi-Step Evaluation Reveals Potential and Bottlenecks</i>	
Masaya Tsunokake, Yuta Koreeda, Terufumi Morishita, Koichi Nagatsuka, Hikaru Tomonari and Yasuhiro Sogawa	304
<i>A Compliance-Preserving Retrieval System for Aircraft MRO Task Search</i>	
Byungho Jo	317
<i>No Label? No Problem: Unsupervised Continual Learning for Adaptive Medical ASR</i>	
Meizhu Liu and Tao Sheng	330
<i>EduPulse: A Practical LLM-Enhanced Opinion Mining System for Vietnamese Student Feedback in Educational Platforms</i>	
Nguyen Xuan Phuc, Phi Nguyen Xuan, Vinh-Tiep Nguyen, Thìn Dang Van and Ngan Luu-Thuy Nguyen	338
<i>When Speed Meets Intelligence: Scalable Conversational NER in an Ever-evolving World</i>	
Karim Ghonim, Antonio Roberto and Davide Bernardi	366
<i>ReflectiveRAG: Rethinking Adaptivity in Retrieval-Augmented Generation</i>	
Akshay Verma, Swapnil Gupta, Siddharth Pillai, Prateek Sircar and Deepak Gupta	377
<i>OCR or Not? Rethinking Document Information Extraction in the MLLMs Era with Real-World Large-Scale Datasets</i>	
Jiyuan Shen, Yuan Peiyue, Atin Ghosh, Yifan Mai and Daniel Dahlmeier	385
<i>PatentVision: A multimodal method for drafting patent applications</i>	
Ruo Yang, Sai Krishna Reddy Mudhiganti and Manali Sharma	397
<i>VideoMind: Thinking in Steps for Long Video Understanding</i>	
Shubhang Bhatnagar, Renxiong Wang, Kapil Krishnakumar, Adel Ahmadyan, Zhaojiang Lin, Lambert Mathias, Xin Luna Dong, Babak Damavandi, Narendra Ahuja and Seungwhan Moon . . .	406
<i>RegNLI: Detecting Online Product Misbranding through Legal and Linguistic Alignment</i>	
Diya Saha, Abhishek Bharadwaj Varanasi, Tirthankar Dasgupta and Manjira Sinha	417

<i>CASPER: Bridging Discrete and Continuous Prompt Optimization through Feedback-Guided Gradient Descent</i>	
Aryan Jain, Pushpendu Ghosh and Promod Yenigalla	425
<i>Adaptive Data Flywheel: Applying MAPE Control Loops to AI Agent Improvement</i>	
Aaditya Shukla, Sidney Knowles, Meenakshi Madugula, David Farris, Ryan Angilly, Santiago Pombo, Lu An, Anbang Xu, Abhinav Balasubramanian, Tan Yu, Jiaxiang Ren and Rama Akkiraju	438
<i>Medical Summarization in Practice: Design, Deployment, and Analysis of a Clinical Summarization System for a German Hospital</i>	
Moiz Rauf and Sean Papay	455
<i>Feedback-Aware Prompt Optimization Framework for Generating Job Postings</i>	
Suraj Maharjan, Ainur Yessenalina and Srinivasan H. Sengamedu	467
<i>Enhancing User Safety: Context-Aware Detection of Offensive Query-Ad Pairs in Multimodal Search Advertising</i>	
Gaurav Kumar, Qiangjian Xi, Tanmaya Shekhar Dabral, Hooshang Ghasemi, Abishek Krishnamoorthy, Danqing Fu, Rui Min, Emilio Antunez, Zhongli Ding and Pradyumna Narayana	475
<i>SAGE: An Agentic Explainer Framework for Interpreting SAE Features in Language Models</i>	
Jiaojiao Han, Wujiang Xu, Mingyu Jin and Mengnan Du	483
<i>Adapting Vision-Language Models for E-commerce Understanding at Scale</i>	
Matteo Nalli, Orshulevich Vladimir, Tala Bazazo, Christian Herold, Michael Kozielski, Marcin Mazur, Szymon Tuzel, Cees G. M. Snoek, Seyyed Hadi Hashemi, Omar Javed, Yannick Versley and Shahram Khadivi	496
<i>MedRiskEval: Medical Risk Evaluation Benchmark of Language Models, On the Importance of User Perspectives in Healthcare Settings</i>	
Jean-Philippe Corbeil, Minseon Kim, Maxime Griot, Sheela Agarwal, Alessandro Sordoni, Francois Beaulieu and Paul Vozila	513
<i>Synthetic Doctor-Patient Dialogue Generation for Robust Medical ASR: A Scalable Pipeline for Vocabulary Expansion and Privacy Preservation</i>	
Kefei Liu and Meizhu Liu	525
<i>Lessons from the Field: An Adaptable Lifecycle Approach to Applied Dialogue Summarization</i>	
Kushal Chawla, Chenyang Zhu, Pengshan Cai, Sangwoo Cho, Scott Novotney, Ayushman Singh, Jonah Lewis, Keasha Safewright, Alfie Samuel, Erin Babinsky, Shi-Xiong Zhang and Sambit Sahu	535
<i>LingVarBench: Benchmarking LLMs on Entity Recognitions and Linguistic Verbalization Patterns in Phone-Call Transcripts</i>	
Seyedali Mohammadi, Manas Paldhe, Amit Chhabra, Youngseo Son and Vishal Seshagiri	545
<i>Improving Training Efficiency and Reducing Maintenance Costs via Language Specific Model Merging</i>	
Alphaeus Dmonte, Vidhi Gupta, Daniel J Perry and Mark Arehart	562
<i>The Subtle Art of Defection: Understanding Uncooperative Behaviors in LLM based Multi-Agent Systems</i>	
Devang Kulshreshtha, Wanyu Du, Raghav Jain, Srikanth Doss, Hang Su, Sandesh Swamy and Yanjun Qi	571
<i>Tailoring Rumor Debunking to You: Diversifying Chinese Rumor-Debunking Passages with an LLM-Driven Simulated Feedback-Enhanced Framework</i>	
Xinle Pang, Danding Wang, Qiang Sheng, Yifan Sun, Beizhe Hu and Juan Cao	586

<i>Synthetic Data Fine-Tuning for Effective Team Formation in Enterprises</i>	
Guilherme Drummond Lima and Adriano Veloso	598
<i>Assertion-Conditioned Compliance: A Provenance-Aware Vulnerability in Multi-Turn Tool-Calling Agents</i>	
Daud Waqas, Aaryamaan Golthi, Erika Hayashida and Huanzhi Mao	610
<i>PROBES : Performance and Relevance Observation for BETter Search</i>	
Sejal Jain, Cyrus Andre DSouza, Jitenkumar Babubhai Rana, Aniket Joshi and Promod Yenigalla	625
<i>Aligning Paralinguistic Understanding and Generation in Speech LLMs via Multi-Task Reinforcement Learning</i>	
Minseok Kim, Jingxiang Chen, Seong-Gyun Leem, Yin Huang, Rashi Rungta, Zhicheng Ouyang, Haibin Wu, Surya Teja Appini, Ankur Bansal, Yang Bai, Yue Liu, Florian Metze, Ahmed A Aly, Anuj Kumar, Ariya Rastrow and Zhaojiang Lin	636
<i>IndicJR: A Judge-Free Benchmark of Jailbreak Robustness in South Asian Languages</i>	
Priyaranjan Pattnayak and Sanchari Chowdhuri	649
<i>Synthesizing question answering data from financial documents: An End-to-End Multi-Agent Approach</i>	
Chetan Harsha, Karmvir Singh Phogat, Sridhar Dasaratha and Shashishekar Ramakrishna	669
<i>Toward Automatic Delegation Extraction in Japanese Law</i>	
Tsuyoshi Fujita, Yuya Sawada, Yusuke Sakai and Taro Watanabe	688
<i>DIALECTIC: A Multi-Agent System for Startup Evaluation</i>	
Jae Yoon Bae, Simon Malberg, Joyce Ann Clarize Galang, Andre Retterath and Georg Groh	711
<i>Long-Context Long-Form Question Answering for Legal Domain</i>	
Anagha Kulkarni, Parin Rajesh Jhaveri, Prasha Shrestha, Yu Tong Han, Reza Amini and Behrouz Madahian	728
<i>ELO: Efficient Layer-Specific Optimization for Continual Pretraining of Multilingual LLMs</i>	
Hangyeol Yoo, ChangSu Choi, Minjun Kim, Seohyun Song, SeungWoo Song, Inho Won, Jongyoul Park, Cheoneum Park and KyungTae Lim	752
<i>MIRAGE: Metadata-guided Image Retrieval and Answer Generation for E-commerce Troubleshooting</i>	
Rishav Sahay, Lavanya Sita Tekumalla and Anoop Saladi	764
<i>CODMAS: A Dialectic Multi-Agent Collaborative Framework for Structured RTL Optimization</i>	
Che-Ming Chang, Prashanth Vijayaraghavan, Ashutosh Jadhav, Charles Mackin, Hsinyu Tsai, Vandana Mukherjee and Ehsan Degan	777
<i>D3: Dynamic Docid Decoding for Multi-Intent Generative Retrieval</i>	
Jaeyoung Kim, Dohyeon Lee, Soona Hong and Seung-won Hwang	789
<i>DisGraph-RP: Graph-Augmented Temporal Modeling with Aspect-Based Contrastive Encoding of Discharge Summary for Readmission Prediction</i>	
Sudeshna Jana, Tirthankar Dasgupta, Manjira Sinha and Pabitra Mitra	801
<i>CareerPathKG: Knowledge Graph Integrated Framework for Career Intelligence</i>	
Ngoc-Quang Le, Duc Duong Hoang, Mai Vu Tran and Thi-Hai-Yen Vuong	813
<i>A Hybrid Supervised-LLM Pipeline for Actionable Suggestion Mining in Unstructured Customer Reviews</i>	
Aakash Trivedi, Aniket Upadhyay, Pratik Narang, Dhruv Kumar and Praveen Kumar	823

<i>ShopperBench: A Benchmark for Personalized Shopping with Persona-Guided Simulation</i>	
Yuan Ling, Chunqing Yuan, Shujing Dong, Yongjian Yang, Nataraj Mocherla and Ayush Goyal	837
<i>ARQA: A Benchmark for Grounded Table–Text QA in Enterprise Annual Reports</i>	
Ruilong Wang and Simone Balloccu	847
<i>Do Clinical Question Answering Systems Really Need Specialised Medical Fine Tuning?</i>	
Sushant Kumar Ray, Gautam Siddharth Kashyap, Sahil Tripathi, Nipun Joshi, Vijay Govindarajan, Rafiq Ali, Jiechao Gao and Usman Naseem	869
<i>SkiLLens: Recognising and Mapping Novel Skills from Millions of Job Ads Across Europe Using Language Models</i>	
Alessia De Santo, Lorenzo Malandri, Fabio Mercorio, Mario Mezzanzanica and Navid Nobani	877
<i>SYMIRECT: A Neuro-Symbolic Divide-Retrieve-Conquer Framework for Enhanced RTL Synthesis and Summarization</i>	
Prashanth Vijayaraghavan, Apoorva Nitsure, Luyao Shi, Charles Mackin, Ashutosh Jadhav, David Beymer, Ehsan Degan and Vandana Mukherjee	886
<i>Benchmarking and Mitigating the Impact of Noisy User Prompts in Medical VLMs via Cross-Modal Reflection</i>	
Zhiyu Xue, Reza Abbasi-Asl and Ramtin Pedarsani	900
<i>Lightweight Domain-Specific Language Model for Real-Time Structuring of Medical Prescriptions</i>	
Jonathan Pattin Cottet, Véronique Eglin and Alex Aussem	915
<i>Balanced Accuracy: The Right Metric for Evaluating LLM Judges - Explained through Youden’s J statistic</i>	
Stephane Collot, Colin Fraser, Justin Zhao, William F. Shen, Timon Willi and Ilias Leontiadis	927
<i>PharmaQA.IT: an Italian dataset for Q&A in the pharmaceutical domain</i>	
Kamyar Zeinalipour, Andrea Zugarini, Asya Zanollo and Leonardo Rigutini	937
<i>DIRECT: Directional Relevance in Conversational Trajectories</i>	
Anshuman Mourya, Rajdeep Mukherjee, Perna Jolly, Vinayak S Puranik and Sivaramakrishnan R Kaveri	948