

ACL 2026

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

Proceedings of the Conference – Volume 6: Industry Track

July 2-7, 2026

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

Diamond

Bloomberg
Engineering




Platinum



Appen



ORACLE



Tencent



北京中关村学院 × 中关村人工智能研究院
Zhongguancun Academy × Zhongguancun Institute of Artificial Intelligence



Gold



JPMorganChase



Silver


Adobe

Bronze



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

ISBN 979-8-89176-394-4

Introduction

We are happy and excited to welcome you to the Industry Track, held under the umbrella of the 64th Annual Meeting of the Association for Computational Linguistics in San Diego, California, USA on 02-07 July 2026.

Like the main research track, the industry track attracted an unprecedented number of submissions: 532 papers! In total, 419 reviewers and 28 area chairs participated in the evaluation of these papers. After a thorough, double-blind peer-review evaluation with three reviews for each submission followed by reviewer discussions and additional deliberations, 153 papers were selected for the ACL 2026 Industry Track, which will be presented as 61 oral talks and 92 posters.

Topic-wise, large language models were once again front and center of almost all submissions with 1. RAG systems and enterprise knowledge AI, 2. agentic systems and workflows, 3. model development, adaptation, optimization, 4. evaluation, benchmarking, LLM assessment, 5. safety, trustworthiness, responsible AI and 6. multimodal, speech and rich content processing being popular topics.

NLP research in academia and NLP research in industry have always been very close in our fields. Our keynote speaker, Roberto Navigli (Sapienza University of Rome and Babelscape, Rome, Italy), will share his insights with regard to this intersection including synergies and fruitful collaborations.

Further insights can be gained through our “Careers in NLP” panel with esteemed participants who have decades of experience with NLP research in academia and industry.

We would like to thank the authors of all Industry Track submissions as well as the reviewers and area chairs for their hard and dedicated work under very tight deadlines. We would also like to thank the General Chair, the Publication Chairs, who supported us in the production of this volume, and all other ACL 2026 committees we interacted with between the summer of 2025, when this endeavour started, and the summer of 2026, when we finally have been able to have the Industry Track at the ACL 2026 conference in San Diego. Finally, we would also like to thank the whole ACL team, especially Jennifer Rachford, as well as the Underline team, especially Damira Mrsic.

Yun Yao Li, Georg Rehm, Mei Tu
Program Co-Chairs

Program Committee

Program Chairs

Yun Yao Li, Adobe
Georg Rehm, Deutsches Forschungszentrum für Künstliche Intelligenz GmbH (DFKI) and Humboldt-Universität zu Berlin
Mei Tu, Samsung

Area Chairs

Amit Agarwal, Oracle
Trung Bui, Adobe Research
Dejing Dou, Fudan University, BEDI Cloud
Gaoyuan Du, Amazon
Xue-Yong Fu, Dialpad Inc.
José Manuel Gómez-Pérez, expert.ai
Leonhard Hennig, Deutsches Forschungszentrum für Künstliche Intelligenz GmbH (DFKI)
Jussi Karlgren, AMD Silo AI and University of Helsinki
Tracy Holloway King, Adobe
Md Tahmid Rahman Laskar, York University and Dialpad Inc.
Jooyoung Lee, Amazon
Pengyuan Li, IBM Research
Xiangfeng Meng, Samsung Research China
Amita Misra, Amazon
Hitesh Laxmichand Patel, Oracle
Priyaranjan Pattanayak, Oracle
Kun Qian, Adobe
Jielin Qiu, Salesforce AI Research
Lisa Raithel, Technische Universität Berlin
German Rigau, Universidad del País Vasco
Natalie Schluter, Apple and IT University
Veronika Solopova, Technische Universität Berlin
Rongguang Wang, Oracle
Chenguang Wang, University of California, Santa Cruz
Chao-Han Huck Yang, NVIDIA Research
Seunghyun Yoon, Adobe Research
Yaping Zhang, Chinese Academy of Sciences
Kai Zhao, Workday AI
Henk van den Heuvel, Radboud University

Reviewers

Asad Abdi, Hadi Abdi Ghavidel, Raia Abu Ahmad, Srinivas Adilapuram, Bhavik Agarwal, Sanjay Agrawal, Zeeshan Ahmed, Mohd Mujtaba Akhtar, Zuhair Akhtar, Sanjeda Akter, Firoj Alam, Jan Alexandersson, Georgios Alexandridis, Kenneth Alperin, Duygu Altinok, Dimitra Anastasiou, Philip Anastassiou, Enrique Henestroza Anguiano, Fnu Antara, Tejas Anvekar, Kushagr Arora, Xabier Arregi, Ankit Arun, Harshita Asnani, Àlex R. Atrio, Sören Auer, Eleftherios Avramidis

Long Bai, Ahsaas Bajaj, Mithun Balakrishna, Anandaganesh Balakrishnan, Ramasubramanian Balasubramanian, Monojit Banerjee, Indervir Singh Banipal, Aditya Bansal, Daniel Bauer, Ian Beaver, Tadesse Destaw Belay, Gabriel Bernier-Colborne, Dario Bertero, Manik Bhandari, Shruti Bhargava, Meghana Moorthy Bhat, Gagan Bhatia, Kasturi Bhattacharjee, Ankita Bhaumik, Nikhil Bhendawade, Daniel M. Bikel, Debmalaya Biswas, Shivangi Bithel, Ekaterina Borisova, Nadjet Bouayad-Agha, Aditya Kiran Brahma, Thomas Brovelli, Yi Bu, Paul Buitelaar, Yash Butala

Ruken Cakici, Dumitru-Clementin Cercel, Gromit Yeuk-Yin Chan, Yashashree Chandak, Nischal Reddy Chandra, Anuj Chauhan, Fuxiang Chen, Guanhua Chen, Jiangning Chen, John Chen, Lin Chen, Wei Chen, Gang Cheng, Zelei Cheng, Zhi-Qi Cheng, Aravind Cheruvu, Purva Chiniya, Shamil Chollampatt, Sanchari Chowdhuri, Ann Clifton, Rylan Conway, Bonaventura Coppola, Jean-Philippe Corbeil, Yuan Cui

Deborah A. Dahl, Daniel Dahlmeier, Daniel Dakota, Souvik Das, Steve DeNeeffe, Daryna Dementieva, Lingjia Deng, Shasvat Mukeshkumar Desai, Shridhar B Devamane, Soumik Dey, Xingjian Diao, Jiayuan Ding, Bin Dong, Andrew Drozdov, Pablo Duboue, Matthew T. Dunn, Sourav Dutta, Jáder Martins Camboim de Sá, Arantza del Pozo

David Elson, Ferhat Erata, Brian Eriksson

Zezhong Fan, Michael Flor, Simona Frenda, Lisheng Fu, Gilad Fuchs, Francesco Fusco

Ramiro H. Gálvez, Swaroop Gadiyaram, Ankur Gandhe, Pengyu Gao, Ankur Garg, Aditya Gautam, Xiou Ge, Anna Lisa Gentile, Piyush Ghai, Diman Ghazi, Reshmi Ghosh, Rikhiya Ghosh, Sayan Ghosh, Filip Ginter, Voula Giouli, Alfio Gliozzo, Jiaying Gong, Naman Goyal, Nidhi Goyal, Yanqi Gu, Boyu Guan, Kalpa Gunaratna, Fu-Ming Guo, Honglei Guo, Tong Guo, Aayush Gupta, Pranav Gupta, Pranjal Gupta, Raghav Gupta, Somya Gupta, Amit Gupta

Amir Hadifar, Udo Hahn, Yuwei Han, Emrul Hasan, Lejian He, Lihong He, Wenchong He, Sanjika Hewavitharana, Derrick Higgins, Arthur Hilbert, Mengze Hong, Pengyu Hong, Wei Hu, Wei-Chieh Huang, Yuchen Hui

Sharon Ibejih, Aditya N Indoori, Samra Irshad, Tomoya Iwakura

Ashish Jain, Gagan Jain, Jitender Jain, Rishabh Jain, Amit Kumar Jaiswal, Miloš Jakubíček, Sai Muralidhar Jayanthi, Shuyue Jia, Li Jin, Rosie Jones

Maithili Sanjay Kadam, Hiroshi Kanayama, Amar Viswanathan Kannan, Vidhyakshaya Kannan, SaiSandeep Kantareddy, Pinar Karagoz, Damianos Karakos, Pradeep Karuturi, Dwipam Katariya, Yannis Katsis, Yoav Katz, Roman Kern, M. Jaleed Khan, Md Abdullah Al Hafiz Khan, Saurabh Khanwalkar, Geewook Kim, Young-Bum Kim, Tracy Holloway King, Miyoung Ko, Sravanth Kodavanti, Svetla Peneva Koeva, George Kour, Jared Kramer, Rajasekar Krishnamurthy, Marek Kubis, Anjishnu Kumar, Priyanshu Kumar, Sachin Kumar, Sanjeev Kumar, Tsung-Ting Kuo

Yanis Labrak, Sarasi Lalithsena, Ian Lane, Stefan Larson, Alexandra Lavrentovich, Grandee Lee, Jochen L. Leidner, Elena Leitner, Changmao Li, Chong Li, Dingcheng Li, Fushan Li, Irene Li, Mingda Li, Pengyuan Li, Sha Li, Xingyu Li, Yingya Li, Yuangang Li, Yupu Liang, Gilbert Lim, Nut Limsopatham, Xiangning Lin, Krister Lindén, Haoran Liu, Jiale Liu, Jingyuan Liu, Pengfei Liu, Ye Liu, Natalia V Loukachevitch, Ziyang Luo

Luyi Ma, Nianzu Ma, Mounica Maddela, Diwakar Mahajan, Akash V Maharaj, Wolfgang Maier,

Fred Mailhot, Ndivhuwo Makondo, Gurshaant Singh Malik, Siddarth Malreddy, Ravi Mandliya, Mudit Mangal, Wenji Mao, Santosh Mashetty, Yuji Matsumoto, John Philip McCrae, David D. McDonald, Alexander Mehler, Tanushree Mehra, Mahnoosh Mehrabani, Jinesh Mehta, Kartik Mehta, Telmo Menezes, Xiangfeng Meng, Margot Mieskes, Hideya Mino, Aditi Mishra, Hemant Misra, Sayantan Mitra, Zhenze Mo, Julian Moreno Schneider, Yasmin Moslem, Lawrence S. Moss, Mahdi Mostajabdaveh, Emir Muñoz, Sidharth Mudgal, Aldrian Obaja Muis, Rohit Mujumdar, Namrata Mukhija, Matthew Mulholland, Ayush Sunil Munot, Syed Shariyar Murtaza, Rithesh Murthy, John Murzaku

Manish Nagireddy, Tetsuji Nakagawa, Jinseok Nam, Diane Napolitano, Krishnasuri Narayanam, Neha, Hoang H Nguyen, Nobal B. Niraula, Behnaz Nojavanasghari

Alexander O'Connor, Oleg Okun, Sergio Oramas, Pedro Ortiz Suarez, Simon Ostermann, Naoki Otani

Stefano Pacifico, Ankur Padia, Manas Paldhe, Alonso Palomino, Bo Pan, Cheoneum Park, Dookun Park, Jeiyoon Park, Seong-Jin Park, Youngja Park, Abhay Dutt Paroha, Patrick Paroubek, Arpankumar Patel, Dhaval C Patel, Sangameshwar Patil, Bibek Paudel, Stephan Peitz, Thang M. Pham, Dhivya Piraviperumal, Jakub Piskorski, Brian Plüss, Lucian Popa, Ravi Teja Potla, Akshara Prabhakar, Pradyot Prakash, Amit Prasad, Shishir Kumar Prasad, Sukannya Purkayastha, Jayeeta Putatunda

Long Qin, Xin Ying Qiu, Elio Querze

Taki Hasan Rafi, Saman Rahbar, Sajjadur Rahman, Aishwarya Rajan, Anshu Rajendra, Poornima Chozhiyath Raman, Nitin Ramrakhiani, Leonardo Ranaldi, Prasanjit Rath, Venkatesh Ravichandran, Traian Rebedea, Ehud Reiter, Yiman Ren, Md Omar Faruk Rokon, Joe Cheri Ross, Weitong Ruan, Mukund Rungta

Harald Sack, Alicia Sagae, Tanmay Sah, Tanay Kumar Saha, Felix Sasaki, Minoru Sasaki, Krati Saxena, Kevin Scaria, David Schlangen, Andreas Schwarz, Robert Schwarzenberg, Rahul Seetharaman, Ethan Selfridge, Shubhashis Sengupta, Ankit Shah, Nirmesh J. Shah, Fu Shang, Mingyue Shang, Yuan Shangguan, Arun Sharma, Raghav Sharma, Sanat Sharma, Michael Shavlovsky, Qiang Sheng, Ashish Shenoy, Kejian Shi, Ibne Farabi Shihab, Tushar Shinde, Vishwajeet Shukla, Chathurangi Shyalika, Andrew Silva, Patrick Simianer, Harvendra Singh, Sneha Singhanian, Raivis Skadiņš, Inguna Skadina, Saleh Soltan, Seonil Son, Hyun-Je Song, Yuanfeng Song, Utkarsh Soni, Pramod Srinivasan, Sebastian Stüker, Katherine Stasaski, Evgeny Stepanov, Dewang Sultania, Chengjie Sun, Suhas Suresha, Shiv Surya, Marek Šuppa

Santosh T.Y.S.S, Narges Tabari, Yuki Tagawa, Mahesh Goud Tandarpally, Yefan Tao, Dung Ngoc Thai, Himanshu Thakur, Vishal Thanvantri Vasudevan, Sudarshan R. Thitte, Philippe Thomas, Shashi Bhushan Tn, Wee Hyong Tok, S.M Towhidul Islam Tonmoy, Manabu Torii, Giuliano Tortoreto, Isabel Trancoso, Keith Trnka, Masaaki Tsuchida

Adrian Ulges, Brian Ulicny, Emmanuel Ngue Um, Max Upravitelev, David Uthus

Praneeth Vadlapati, Tom Vanallemeersch, Andrea Varga, Vaibhav Varshney, Manisha Verma, Nikhil Verma, Martin Villalba, Emmanuel Vincent, Rahul Deo Vishwakarma

Kapil Wanaskar, Bingqing Wang, Chengyu Wang, Hai Wang, Jin Wang, Mingxian Wang, Ryan Wang, Yi Wang, Yile Wang, Yuying Wang, Zhengxiang Wang, Azmine Toushik Wasi, Penghui

Wei, Haryo Akbarianto Wibowo, Ben Wiesel, Andreas Witt, Jinting Wu, Meng-Chen Wu, Tianxing Wu

Jing Xia, Lu Xiang, Tong Xiao, Kaige Xie, Zejun Xie, Guangyue Xu, Jun Xu

Victor Yang, Wen Yang, Yeongwook Yang, Jinyeong Yim, Yuwei Yin, Lei Yu, Mingzhi Yu

Fadi Zaraket, Chen Zhang, Dan Zhang, Fan Zhang, Lei Zhang, Mengchao Zhang, Tianlin Zhang, Wenbo Zhang, Yin Zhang, Yudi Zhang, Yunbai Zhang, Dong Zhou, Yingxue Zhou, Junnan Zhu, Xiliang Zhu, Xuan Zhu, Yimeng Zhuang, Wenjie Zi

Table of Contents

<i>ACL 2026 Industry Track: Overview</i>	
Yunyao Li, Georg Rehm and Mei Tu	1
<i>Smarter, not Bigger: Fine-Tuned RAG-Enhanced LLMs for Automotive Hardware-in-the-Loop Testing</i>	
Chao Feng, Zihan Liu, Siddhant Gupta and Jan von der Assen	5
<i>MinerU2.5: A Decoupled Vision-Language Model for Efficient High-Resolution Document Parsing</i>	
Junbo Niu, Zheng Liu, Zhuangcheng Gu, Bin Wang, Linke Ouyang, Zhiyuan Zhao, Tao Chu, Tianyao He, Fan Wu, Qintong Zhang, Zhenjiang Jin, Guang Liang, Rui Zhang, Wenzheng Zhang, Yuan Qu, Zhifei Ren, Yuefeng Sun, Zirui Tang, Boyu Niu, Yuanhong Zheng, Dongsheng Ma, Ziyang Miao, Hejun Dong, Siyi Qian, Junyuan Zhang, Fangdong Wang, Jingzhou Chen, Xiaomeng Zhao, Liqun Wei, Wei Li, Shasha Wang, RuiLiang Xu, Yuanyuan Cao, Lu Chen, Qianqian Wu, Huaiyu Gu, Lindong Lu, Dechen Lin, Shenguanlin, Xuanhe Zhou, Linfeng Zhang, Yuhang Zang, Xiaoyi Dong, Jiaqi Wang, Bo Zhang, Lei Bai, Pei Chu, Weijia Li, Jiang Wu, Lijun Wu, Zhenxiang Li, Guangyu Wang, Zhongying Tu, Chao Xu, Kai Chen, Bowen Zhou, Dahua Lin, Wentao Zhang and Conghui He	13
<i>Leveraging Generative AI for Extracting Business Requirements from Legacy COBOL and PL/I Code</i>	
Ankur Kalohia	43
<i>Is Agentic RAG worth it? An experimental comparison of RAG approaches</i>	
Pietro Ferrazzi, Milica Cvjetićanin, Alessio Piraccini and davide giannuzzi	56
<i>Improving Hate Speech Detection by Fusing Textual and User Interaction Representations in Online Communities</i>	
Xu Gao, Dong Jing and Kee-hung Lai	76
<i>TeachMaster: Generative Teaching via Code</i>	
Yuheng Wang, Runde Yang, Lin Wu, Jie Zhang, Jingru Fan, Tianle Zhou, Ruoyu Fu, Huatao Li, Ruijie Shi, Siheng Chen, Weinan E and Chen Qian	88
<i>ARGUS: Policy-Adaptive Ad Governance via Evolving Reinforcement with Adversarial Umpiring</i>	
Deyi Ji, Junyu Lu, Xuanyi Liu, Liqun Liu, Hailong Zhang, Peng Shu, Huan Yu, Jie Jiang, Tianrun Chen and Lanyun Zhu	99
<i>Synthetic Text Detection in the Age of Large Language Models: Watermark vs. Automatic Detection</i>	
Adaku Uchendu	113
<i>R³A: Reinforced Reasoning for Relevance Assessment for RAG in User-Generated Content Platforms</i>	
Xiaowei Yuan, Lei Jin, Haoxin Zhang, Ziyang Huang, Yan Gao, Yiwu, Yao Hu, Jun Zhao and Kang Liu	127
<i>HAD: HALLucination Detection Language Models Based on a Comprehensive Hallucination Taxonomy</i>	
Fan Xu, Xinyu Hu, Zhenghan Yu, Li Lin, Xu Zhang, Yang Zhang, Wei Zhou, Jinjie Gu and Xiaojun Wan	141
<i>DocSplit: A Comprehensive Benchmark Dataset and Evaluation Approach for Document Packet Recognition and Splitting</i>	
Md Mofijul Islam, Md Sirajus Salekin, Nivedha Balakrishnan, Vincil C. Bishop III, Niharika Jain, Spencer Romo, Bob Strahan, Boyi Xie and Diego A. Socolinsky	159
<i>TaoType: Predicting Fine-Grained Typing Intent for Faster Search</i>	
Yipeng Yu, Yichen Yuan, Chengxiao Feng and Xu Liu	184

<i>Calibrating LLM Judges: Linear Probes for Fast and Reliable Uncertainty Estimation</i>	
Bhaktipriya Radharapu, Eshika Saxena, Kenneth Li, Chenxi Whitehouse, Adina Williams and Nicola Cancedda	193
<i>Detecting AI-Generated Content on Social Media with Multi-modal Language Models</i>	
Chenyang Yang, Shen Yan, Yibo Yang, Litao Hu, Yuchen Liu, Yuan Zeng, Hanchao Yu, Yinan Zhu, Sumedha Singla, Brian Vanover, Huijun Qian, Zihao Wang, Fujun Liu, Aashu Singh, Jianyu Wang and Xuewen Zhang	217
<i>IceBreaker for Conversational Agents: Breaking the First-Message Barrier with Personalized Starters</i>	
Hongwei Zheng, Weiqi Wu, Zhengjia Wang, Guanyu Jiang, Haoming Li, Tianyu Wu, Yongchun Zhu, Jingwu Chen and Feng Zhang	230
<i>Motivating Next-Gen Accelerators with Flexible N:M Activation Sparsity via Benchmarking Lightweight Post-Training Sparsification Approaches</i>	
Shirin Alanova, Kristina Kazistova, Ekaterina Galaeva, Alina Kostromina, Vladimir Smirnov, Redko Dmitry, Alexey Dontsov, Maxim Zhelnin, Evgeny Burnaev and Egor Shvetsov	242
<i>Sommelier: Scalable Open Multi-turn Audio Pre-processing for Full-duplex Speech Language Models</i>	
Kyudan Jung, Jihwan Kim, Soyeon Kim, Jeonghoon Kim, Jaegul Choo and Cheonbok Park .	259
<i>Proximity-Based Multi-Turn Optimization: Practical Credit Assignment for LLM Agent Training</i>	
Yangyi Fang, Jiaye Lin, Xiaoliang Fu, Cong Qin, Haolin Shi and Chang Liu	285
<i>Mixture of Heterogeneous Grouped Experts for Language Modeling</i>	
Zhicheng Ma, Xiang Liu, Zhaoxiang Liu, Ning Wang, Yi Shen, Kai Wang, Shuming Shi and Shiguo Lian	308
<i>Policy Compliance of User Requests in Natural Language for AI Systems</i>	
Pedro Cisneros-Velarde	318
<i>Better Literary Translation: A Multi-Aspect Data Generation and LLM Training Approach</i>	
Zhihao Lin, Ziqi Zhu, Hao Huang, Guanghui Wang and Peiyang He	329
<i>SSR-A: Spatial- and Semantic-Aware Instructions and Curriculum Reinforcement for Advertisement Compliant Rectification</i>	
Cao Te, Mengge Xue, zhenyu hu, Yuan Chen, Liqun Liu, Peng Shu, Huan Yu and Jie Jiang . .	346
<i>\mathcal{R}^3: Advertisement Compliance Rectification via Group-Relative Experience Extractor and Curriculum Reinforcement</i>	
Yuan Chen, zhenyu hu, Mengge Xue, Cao Te, Liqun Liu, Peng Shu, Huan Yu and Jie Jiang . .	358
<i>GroCLM: Grocery Category Recommendation in E-Commerce with Large Language Models</i>	
Yuan Zhong, Chuanwei Ruan, Moein Hasani, Tejaswi Tenneti, Haixun Wang and Fenglong Ma	371
<i>DeepResearch Retail: Benchmarking Tool-Augmented Deep Research in the E-Commerce Domain</i>	
Rafael Ferreira, Flavio Di Palo, Huilin Lu, Ayush Jain and Harsha Aduri	386
<i>GeoGround: Uncertainty-Weighted Multi-Task Learning for Geo-Alignment and Address Defect Detection</i>	
Srinivas Virinchi, Aman Gulati and Anoop Saladi	410
<i>Pre-Deployment Advertisement Ranking under Data Scarcity via Context-Aware Criteria Generation with VLMs</i>	
Kyunggho Kim, Yeonje Choi, Gyurim Hwang, Sejin Chung, Hongseok Lee, Myeong Ho Song, Yeonggho Kim, Sunwoo Kim, Jongha Lee, Juyeon Kim and Kijung Shin	420

<i>Agent-Ops: A Multi-Agent Orchestration Framework for End-to-End SOP Automation in E-Commerce Operations</i>	
Apoorva Singh, Sanjay Agrawal, Sayanta Adhikari, Vinayak S Puranik, Shivam Tiwari and Dheer-aj Assudani	436
<i>Reason-Code: Reliable Code Generation via Test-Driven Monte Carlo Tree Search</i>	
Zixu Li and Zhiqi Peng	447
<i>Optimizing Retrieval-Augmented Generation for E-Commerce How-To Assistance</i>	
Gilad Fuchs, Leonid Ekimov, Fei Dong, Jiahong Xie, Wei Liu, Maxim Manco and Alexander Nus	459
<i>UNIVID: Unified Vision-Language Model for Video Moderation</i>	
Kejuan Yang, Yizhuo Zhang, Mingyuan Du, Yue Zhang, Dixin Zheng, Kaili Zhao, Yang Xiao, Hanzhong Liang and Kenan Xiao	467
<i>Learning from Textual Radiology Reports: A Benchmark Dataset for Coronary CT Angiography</i>	
Sudharshan Balaji, Zhiyu Liu, Zhengyuan Jiang, Shuo Lei, Yimin Chen, Yang Xiao, Shone O. Almeida, Mathew Joseph Karivelil, Christopher Malanga and Ning Wang	480
<i>Teaching LLM to be Persuasive: Reward-Enhanced Policy Optimization for Alignment from Heterogeneous Rewards</i>	
Xia Zeng, Yihan Chen, Luhui Liu, Chao Luo, Ye Chen and Zhuangzhuoran	494
<i>NEST: Nested Evidence Survival for Retrieval</i>	
Akshay Verma, Siddharth Pillai, Prateek Sircar and Deepak Gupta	507
<i>Search-P1: Path-Centric Reward Shaping for Stable and Efficient Agentic RAG Training</i>	
Tianle Xia, Ming Xu, Lingxiang Hu, Yiding Sun, wenwei li, Linfang Shang, Liqun Liu, Peng Shu, Huan Yu and Jie Jiang	517
<i>AgenticQwen: Training Small Agentic Language Models with Dual Data Flywheels for Industrial-Scale Tool Use</i>	
Yuanjie Lyu, Chengyu Wang, Haonan Zheng, Yuanhao Yue, Junbing Yan, Ming Wang and Jun Huang	535
<i>From Short Video to Clickable Search: RLVR-Enabled Listwise Query Suggestion with Retrieval-Augmented Context</i>	
Mingkai Tian, xuye, Long Meng, Liwei Chen, Zhiheng Qin and Yi Wang	552
<i>Small Agents, Big Gains: Journey-Aware and Critic-Guided Simulation for Long-Horizon Shopping Dialogues</i>	
Qing Ping, Changyou Chen and Binxuan Huang	563
<i>Progressive Fine-Tuning for Cost-Effective Structured Attribute Generation in E-commerce</i>	
Lakshman Kolasani and Fatemeh Taheri Dezaki	585
<i>Grounded Multimodal In-Context Learning for Product Weight Estimation at Scale in E-commerce</i>	
Bhavuk Singhal, Arsh Keshari and Ravindra Kumar Yadav	592
<i>Towards Faithful Industrial RAG: A Reinforced Co-adaptation Framework for Advertising QA</i>	
wenwei li, Ming Xu, Tianle Xia, Lingxiang Hu, Yiding Sun, Linfang Shang, Liqun Liu, Peng Shu, Huan Yu and Jie Jiang	604
<i>A Unified Framework for Modeling Heterogeneous Financial Data via Dual-Granularity Prompting</i>	
Yu Lei, zixuan wang, Yiqing Feng, Junru Zhang, Yahui Li, Liu Chu, Wang Tongyao and Dongyang Li	618

<i>DuIVRS-2: An LLM-based Interactive Voice Response System for Large-scale POI Attribute Acquisition</i>	
Le Zhang, Shengming Zhang, Rui Zha, Wu Yunpeng, Jingbo Zhou and Jizhou Huang	632
<i>SAJA: A Simple Approach to Judge Alignment for LLM-as-a-Judge</i>	
Sneha Kola, Pankaj Kumar Sharma, Soumyadeep Dey, Bamdev Mishra and Mayur Datar	646
<i>ColorBrowserAgent: Complex Long-Horizon Browser Agent with Adaptive Knowledge Evolution</i>	
Jihong Wang, Jiamu Zhou, Weiming Zhang, Teng Wang, Weiwen Liu, Zhuosheng Zhang, Xingyu Lou, Weinan Zhang, Huarong Deng and Jun Wang	665
<i>ClinDet-Bench: Beyond Abstention, Evaluating Judgment Determinability of LLMs in Clinical Decision-Making</i>	
Yusuke Watanabe, Yohei Kobashi, Takeshi Kojima, Yusuke Iwasawa, Yasushi Okuno and Yutaka Matsuo	681
<i>AIDA-SEAT: Towards Reliable AI Doctor Assistant via State-Evaluation-Action Tree Enhanced LLMs in Online Hospital</i>	
Lianxin Sun, Xiaoying Ying, Guangya Yu, Weiyan Zhang, Chenhao Guan, Hao He, Mingxi Shang, Jianhua Li, ChunMing Wang and Tong Ruan	704
<i>IndustryAssetEQA: A Neurosymbolic Operational Intelligence System for Embodied Question Answering in Industrial Asset Maintenance</i>	
Chathurangi Shyalika, Dhaval C Patel and Amit Sheth	716
<i>Evidence-Driven Reasoning for Industrial Maintenance Using Heterogeneous Data</i>	
Fearghal O'Donncha, Nianjun Zhou, Natalia Martinez, James T Rayfield, Fenno F. Heath III, Abigail Langbridge and Roman Vaculin	736
<i>MobileLLM-Flash: Latency-Guided On-Device LLM Design for Industry Scale Deployment</i>	
Hanxian Huang, Igor Fedorov, Andrey Gromov, Bernard Beckerman, Naveen Suda, David Eriksson, Maximilian Balandat, Rylan Conway, Patrick Huber, Chinnadhurai Sankar, Ayushi Dalmia, Zechun Liu, Lemeng Wu, Tarek Elgamal, Adithya Sagar, Vikas Chandra and Raghuraman Krishnamoorthi	749
<i>LEAP: Layer-wise Exit-Aware Pretraining for Efficient Transformer Inference</i>	
Shashank Kapadia, Deep Narayan Mishra, Sujal Reddy Alugubelli, Haoan Wang, Saipraveen Vabbilisetty, Rishi Bhatia and Anupriya Sharma	761
<i>Think Like You Execute: Verifiable Chain of Thought from Program Traces</i>	
Shailja Thakur, Vaibhav Saxena, Rohan Kulkarni, Shivdeep Singh, Parameswaran Selvam, Hiroshi Kanayama and Hima Patel	775
<i>From Relevance to Authority: Authority-aware Generative Retrieval in Web Search Engines</i>	
Sunkyung Lee, Jihye Back, Donghyeon Jeon, Soonhwan Kwon, Moonkwon Kim, Inho Kang and Jongwuk Lee	796
<i>Enhancing Online Recruitment with Category-Aware MoE and LLM-based Data Augmentation</i>	
Minping Chen, Bing Xu, Zulong Chen, Chuanfei Xu, Ying Zhou, Zui Tao and Zeyi Wen	812
<i>EfficientTool: A Cost-Effective Aligning Framework for Tool-Conditioned Agents in SME Scenarios</i>	
Yuanqi Mu, Bingfeng.Pi, Defei Xia, Lei.Zuo and Yongqi Zhang	825
<i>Unleashing Low-Bit Inference on Ascend NPUs: A Comprehensive Evaluation of HiFloat Formats</i>	
Pengxiang Zhao, Hui-Ling Zhen, Xing Li, Han Bao, Weizhe Lin, Zhiyuan Yang, Yu Zi Wei, Xin Wang, Mingxuan Yuan, Xianzhi Yu and Zhenhua Dong	844

<i>Securing the Tool Layer: A Threat Taxonomy and Runtime Defense Framework for Model Context Protocol Deployments</i>	
Saurabh Yergattikar	865
<i>Enhancing Job Evaluation with Data Augmentation and Text Classification</i>	
Samaneh Jalilian, Niels van Weeren, Mohammad Shokri, Thijmen Bijl and Suzan Verberne .	872
<i>When Good OCR Is Not Enough: Benchmarking OCR Robustness for Retrieval-Augmented Generation</i>	
Lin Sun, wangdexian, Jingang Huang, Linglin Zhang, Change Jia, Zhengwei Cheng and Xi-angzheng Zhang	884
<i>Aggregate vs. Personalized Judges in Business Idea Evaluation: Evidence from Expert Disagreement</i>	
Wataru Hirota, Tomoki Taniguchi, Tomoko Ohkuma, Kosuke Takahashi, Takahiro Omi, Kosuke Arima, Takuto Asakura, Chung-Chi Chen and Tatsuya Ishigaki	895
<i>Memory-Efficient Structured Backpropagation for On-Device LLM Fine-Tuning</i>	
JuneYoung Park, Yuri Hong, Seongwan Kim and Jaeho Lee	906
<i>Multi-Agent Orchestration for Terminology-Constrained Machine Translation in Industrial Localiza-tion</i>	
Emanuele Di Rosa	917
<i>TrendPulse: A Simple yet Efficient Framework for Capturing Viral E-Commerce Spikes via LLM-Driven Contextualization</i>	
Arin Jain, Devashish Gupta, Bhavuk Singhal, Divay Jindal, Vinit Rongata and Ravindra Kumar Yadav	927
<i>ACRM: Multi-Agent Trajectory Learning for Automated Credit Risk Model Refreshing in Production</i>	
Liangzu Liu, Mengzhe Ruan, Xiaotian Chen, HaonanChen, XudongNiu, Wendi Yuan, YuechenLi, Yang Liu and GuanJun Wang	942
<i>From 2:4 to 8:16 sparsity patterns in LLMs for Outliers and Weights with Variance Correction</i>	
Egor Maximov, Yulia Kuzkina, Egor Shvetsov, Azamat Kanametov, Aleksandr Prutko, Maxim Zhelnin and Aleksei Goncharov	957
<i>An Efficient Framework for Whole-Page Reranking via Single-Modal Supervision</i>	
Zishuai Zhang, Sihao Yu, Xiewenyi, Ying Nie, Junfeng Wang, Zhiming Zheng, Dawei Yin and Hainan Zhang	966
<i>LLM as a Risk Manager: LLM Semantic Filtering for Lead-Lag Trading in Prediction Markets</i>	
Sumin Kim, Minjae Kim, Jihoon Kwon, Yoon Kim, Oscar Levy, Alejandro Lopez-Lira, Yongjae Lee and Chanyeol Choi	979
<i>Diffusion-Pretrained Dense and Contextual Embeddings</i>	
Sedigheh Eslami, Maksim Gaiduk, Markus Krimmel, Louis Mark Milliken, Bo Wang and Denis Bykov	990
<i>Actionable Interpretability for Churn Classification: A Text Bottleneck Model Case Study at a Major Telecom Provider</i>	
Adrian Sauter, Vera Neplenbroek, Georgios Vlassopoulos and Gianluigi Bardelloni	1005
<i>Decoding Text Spans for Efficient and Accurate Named-Entity Recognition</i>	
Andrea Maracani, Savas Ozkan, Junyi Zhu, Sinan Mutlu and Mete Ozay	1025
<i>EVE: A Domain-Specific LLM Framework for Earth Intelligence</i>	
Alex R. Atrio, Antonio Lopez, Jino Rohit, Yassine El Ouahidi, Marcello Politi, Vijayasri Iyer, Umar Jamil, Sébastien Bratières and Nicolas Longépé	1037

<i>Hemolix.TabGen: Optimized Table Generation from Documents</i>	
Gyanendra Shrestha, Todor Ivanov, Karthik Vemireddy, Anna Pyayt and Michael Gubanov	1055
<i>Harmonizing Dense and Sparse Signals in Multi-turn RL: Dual-Horizon Credit Assignment for Industrial Sales Agents</i>	
Haojin Yang, Ai Jian, Yiwei Wang, Xinyue Huang, Weipeng Zhang, Ke Zeng, Xunliang Cai and Jingqing Ruan	1067
<i>Know What You See: Grounded localization of product components</i>	
Manan Soni, Abinesh Kanagarajan and Shyam Mohan	1077
<i>SCAIR: Schema-Conditioned Agentic Iterative Reasoning for Enterprise Knowledge Graphs</i>	
Prateek Chaturvedi, Yuqicheng Zhu, Hongkuan Zhou, Dongzhuoran Zhou, Yunjie He, Steffen Staab, Fei Du, Jie Tang and Evgeny Kharlamov	1089
<i>FROST: Factual Reasoning via Optimized Stochastic Trajectories in Large Language Models during Inference</i>	
Soumedhik Bharati, Ebad Shabbir and Jiechao Gao	1105
<i>SchemaRAG: Dynamic Large Schema Reduction for LLM-driven Structured Information Extraction</i>	
Sin Yu Bonnie Ho, Arlie Coles, Erik Larsson, Eric Marshall, Nathan Bodenshtab and Paul Vozila	1114
<i>DQA: Diagnostic Question Answering for IT Support</i>	
Vishaal Kapoor, Mariam Dundua, Evren Yortucboyulu, Sarthak Ahuja, Neda Kordjazi, Yiming Li, Vaibhavi padala, Derek Ho, Jennifer Whitted and Rebecca Steinert	1128
<i>Agentic Economic Modeling</i>	
Bohan Zhang, Jiaxuan Li, Ali Hortacsu, Xiaoyang Ye, Victor Chernozhukov, Anqi Ni and Edward W Huang	1136
<i>Robust Explanations for User Trust in Enterprise NLP Systems</i>	
Guilin Zhang, Kai Zhao, Jeffrey Friedman, Xu Chu and Amine Anoun	1150
<i>VEG: Verbal ϵ-greedy for Semantic Exploration in Multi-Turn RL Agents</i>	
Yongchang Hao, Jie Hao, Yongsheng Mei, Ze Ye, Junyi Chai, Bin Guo, Benjamin Z. Yao, Chenlei Guo and Lili Mou	1159
<i>FlowHN: Adaptive Token Routing for Efficient Parallel Hybrid Networks</i>	
Mohammad Mahdi Moradi, Walid Ahmed, Shuangyue Wen, Sudhir Mudur, Weiwei Zhang and Yang Liu	1170
<i>MemoPhishAgent: Memory-Augmented Multi-Modal LLM Agent for Phishing URL Detection</i>	
Xuan Chen, Hao Liu, Tao Yuan, Mehran Kafai, Piotr Habas and Xiangyu Zhang	1182
<i>Adaptive Weighted Proxy Tuning: Efficient Gray-Box Steering for Image Captioning.</i>	
Nafew Azim, Fuad Rahman and Nabeel Mohammed	1197
<i>CLARITY: A Framework and Benchmark for Conversational Language Ambiguity and Unanswerability in Interactive NL2SQL Systems</i>	
Tabinda Sarwar, Farhad Moghimifar, Cong Duy Vu Hoang, Xiaoxiao Ma, Shawn Chang Xu, Fahimeh Saleh, Poorya Zaremoondi, Avirup Sil and Katrin Kirchhoff	1218
<i>Toward Scalable Verifiable Reward: Proxy State-Based Evaluation for Multi-turn Tool-Calling LLM Agents</i>	
Yun-Shiuan Chuang, Chaitanya Kulkarni, Alec M. Chiu, Avinash Thangali, Zijie Pan, Shivani Shekhar, Yirou Ge, Yixi Li, Uma Kona, Linsey Pang and Prakhar Mehrotra	1251

<i>Ada-RS: Adaptive Rejection Sampling for Selective Thinking</i>	
Yirou Ge, Yixi Li, Alec M. Chiu, Shivani Shekhar, Zijie Pan, Avinash Thangali, Yun-Shiuan Chuang, Chaitanya Kulkarni, Uma Kona, Linsey Pang and Prakhar Mehrotra	1265
<i>IPS: In-Prompt Process Supervision for Short Video Content Moderation</i>	
Mingchao Liu, Yu Sun, Ruixiao Sun, Xin Dong, Xiang Shen, Hongwei Wang, Hongyu Xiong and Yang Song	1277
<i>RedOne 2.0: Rethinking Domain-specific LLM Post-Training in Social Networking Services</i>	
Fei zhao, Chonggang Lu, Haofu Qian, Fangcheng Shi, Zijie Meng, Jianzhao Huang, Zheyong Xie and Shaosheng Cao	1289
<i>How to Train a Real-World Silicon Concierge? Internalizing Complex Business Workflow to Only OneModel</i>	
Yongqi Tong, Xiaoyun Feng, Lyuxin Xue, Jianshe Li, Xin Zhang and Jiang-Ming Yang	1303
<i>FinRAG-12B: A Production-Validated Recipe for Grounded Question Answering in Banking</i>	
Denys Katerenchuk, Pablo Duboue and Keelan Evanini	1317
<i>CascadeDebate: Multi-Agent Deliberation for Cost-Aware LLM Cascades</i>	
Raeyoung Chang, Dongwook Kwon, Jisoo Lee and Nikhil Verma	1329
<i>SemViQA: A Semantic Question Answering System for Vietnamese Information Fact-Checking</i>	
Dien X. Tran, Nam V. Nguyen, Tran Tan Thanh, Anh T. Hoang, Dng Văn Tài, Lê Thanh Di and Phuc-Lu Le	1341
<i>SimRPD: Optimizing Recruitment Proactive Dialogue Agents through Simulator-Based Data Evaluation and Selection</i>	
Zhiyong Cao, Dunqiang Liu, Qi Dai, Haojun Xu, Huai Yuen Khor, Hao Wang, Huan He, Yafei Liu, Ke Ma, Ruqian Shi, Sicheng Zhou and Sijia Yao	1359
<i>ProductResearch: Training E-Commerce Deep Research Agents via Multi-Agent Synthetic Trajectory Distillation</i>	
Jiangyuan Wang, Kejun Xiao, Huaipeng Zhao, Tao Luo and Xiaoyi Zeng	1378
<i>ReList: A Multi-objective Reasoning Framework for Diversified Listwise Query Recommendation</i>	
Shuxian Bi, Chenxu Wang, Wenjie Wang, Yueqi Mou, Fuli Feng, Tang Biao and Peng Yan	1392
<i>Optimal Expert-Attention Allocation in Mixture-of-Experts: A Scalable Law for Dynamic Model Design</i>	
Junzhuo Li, Peijie Jiang, Changxin Tian, Jia Liu, Zhiqiang Zhang and Xuming Hu	1406
<i>A Multistage Extraction Pipeline for Long Scanned Financial Documents: An Empirical Study in Industrial KYC Workflows</i>	
Han Yuxuan, Yuanxing Zhang, Yushuo Wang and Yichao Jin	1419
<i>Effective Performance Measurement: Challenges and Opportunities in KPI Extraction from Earnings Calls</i>	
Rasmus T. Aavang, Rasmus Tjalk-Bøggild, Alexandre Iolov, Giovanni Rizzi, Mike Zhang and Johannes Bjerva	1434
<i>Structure-Guided Entity Resolution: Fine-Tuning LLMs for Robust Name Matching in Complex Linguistic Contexts</i>	
Shivam Chourasia, Hitesh Kapoor and Nilesh Patil	1461
<i>Accurate Legal Reasoning at Scale: Neuro-Symbolic Offloading and Structural Auditability for Robust Legal Adjudication</i>	
Stanisław Sójka and Witold Kowalczyk	1469

<i>CI-Work: Benchmarking Contextual Integrity in Enterprise LLM Agents</i>	
Wenjie Fu, Xiaoting Qin, Jue Zhang, Qingwei Lin, Lukas Wutschitz, Robert Sim, Saravan Rajmohan and Dongmei Zhang	1483
<i>From Pixels to Policies: Reinforcing Spatial Reasoning in Language Models for Content-Aware Layout Design</i>	
Sha Li, Stefano Petrangeli, Yu Shen and Xiang Chen	1509
<i>Practical Guidelines for Model Merging in LLMs Pre-Training</i>	
Giuseppe Curci, Stefano Simonazzi, Andrea molinari and Andrea Zugarini	1519
<i>Are Large Language Models Economically Viable for Industry Deployment?</i>	
Abdullah Mohammad, Sushant Kumar Ray, Pushkar Arora, Rafiq Ali, Ebad Shabbir, Gautam Siddharth Kashyap, Jiechao Gao and Usman Naseem	1533
<i>Let LLM Tutors Ask First: Proactive LLM-Based Tutoring at Scale in a 1,500-Student Online Classroom</i>	
Jonghoon Lee, Geonjae Youn, Seongmin Lee, Chaemoon Im, Joongheon Kim and Chuck Yoo	1541
<i>Hierarchical Reason-of-Contact Detection in Retail Banking Customer Interactions via LLM-Driven Taxonomy Induction</i>	
Yoages Kumar Mantri, Yagneswaran Sriraja, John Hartley, Greig A Cowan, Raad Khraishi and Abhinav Gupta	1555
<i>Fanar-Sadiq: A Multi-Agent Architecture for Grounded Islamic QA</i>	
Ummar Abbas, Mourad Ouzzani, Mohamed Y. Eltabakh, Omar Sinan, Gagan Bhatia, Hamdy Mubarak, Majd Hawasly, Mohammed Qusay Hashim, Kareem Mohamed Darwish and Firoj Alam	1567
<i>GuideTree: Guideline-Induced Review Trees for Long Medical Records</i>	
Chengze Ge, Ruiqing Zhang, Yining Wang, Shengping Liu, Liang Jiaen, Weihuang and Weihua Chen	1589
<i>Automatic Prompt Engineering for Scalable Prompt Inversion in Text-to-Image Ad Generation</i>	
Zixin Ding, Qi Zeng, Boying Gong, Wenlong Deng, Bo Pan and Yuxin Chen	1605
<i>Efficient Agent Evaluation via Diversity-Guided User Simulation</i>	
Itay Nakash, George Kour and Ateret Anaby Tavor	1627
<i>What Question Did You Answer? Refining Contact Center Evaluation Plans via Backward Questions</i>	
Prajwal Sood, Rushikesh Pawar, Digvijay Anil Ingle and Anup Pattnaik	1649
<i>Natural-Language Policies to Executable Decisions: An Interpretable Large Language Model Framework</i>	
Ziqiang Zhang, Jing Ma, Zilong Wang, Jiayuan Chen, Yi Qiao, Yu He, Wei Zhang, Dai Cheng and Xiaoyu Shen	1668
<i>Diagnose, Then Repair: A Two-Stage MQM-Guided Post-Editing Framework for Domain-Specific Machine Translation</i>	
Ji Hun Wang and Siyu Wu	1683
<i>Don't Stop Early: Scalable Enterprise Deep Research with Controlled Information Flow and Evidence-Aware Termination</i>	
Prafulla Kumar Choubey, Kung-Hsiang Huang, Pranav Narayanan Venkit, Jiaxin Zhang, Vaibhav Vats, Yu Li, Xiangyu Peng and Chien-Sheng Wu	1699

<i>FinHarmBench: Financial Jailbreak Benchmark and Unsupervised Safety Fine-Tuning via Refusal Steering Distillation</i>	
Yubin Choi, Yujin Yang, Subin Kim, Seokil Ham, Seungju Cho, Jungmin Son, Youngjun Kwak and Changick Kim	1714
<i>PAR²-RAG: Planned Active Retrieval and Reasoning for Multi-Hop Question Answering</i>	
Xingyu Li, Rongguang Wang, Yuying Wang, Mengqing Guo, Chenyang Li, Tao Sheng, Sujith Ravi and Dan Roth	1727
<i>OmniOData: Unleashing Small Language Models for OData Query Generation with Synthetic Data and Reinforcement Learning</i>	
Tao Bai, Zhaochen Li, Hongxin Shao and Daniel Dahlmeier	1738
<i>LegalDrill: Diagnosis-Driven Synthesis for Legal Reasoning in Small Language Models</i>	
Tianchun Li, Haochen Liu, Vishwa Pardeshi, Xingchen Wang, Tianci Liu, Huijun Zhao, Wei Fan and Jing Gao	1755
<i>From TextBlob to LLM Agents: Sentiment Model Selection for B2B Technical Support with CSAT Ground Truth</i>	
Pedro Vidigal	1774
<i>A Novel Matching Paradigm: Unified Generative and Discriminative LLM with Prompt Compression for Relevance Learning</i>	
Guoliang Zhao, Zixin Cui, Chao Ye, Dengwu He, Fei Huang, Yubo Liu, shuanglong li, Tzungren Kuo, Bin Ding, Shuang Zhang, KunhongZhu, Zhi Guo and Liu Lin	1783
<i>Rethinking Scale: Deployment Trade-offs of Small Language Models under Agent Paradigms</i>	
Xinlin Wang	1795
<i>FourCorners: A Production Knowledge Graph Unifying Thailand's Legal System</i>	
Pawitsapak Akarajaradwong, Sarana Nutanong and Chompakorn Chaksangchaichot	1808
<i>Modular Monolingual Adaptation using Pretrained Language Models</i>	
Nalin Kumar and Ondrej Dusek	1819
<i>Multi-Task Reinforcement Learning for Enhanced Multimodal LLM-as-a-Judge</i>	
Junjie Wu, Xuan Kan, Zihao He, Shunwen Tan, Bo Pan and Kaitai Zhang	1829
<i>RAPIDS: Resume Attack Prompt Injection Detection at Scale</i>	
Yohann Augey, Joshua H. Levy and Arda Akdemir	1848
<i>No Innocence in Styling: Discovery of Privacy Protection Capabilities and Security Risks in Consumer Generative AI Writing Assistants</i>	
Mohd. Farhan Israk Soumik, Syed Mhamudul Hasan, Wanniarachchi Kankanamge Malithi Mithsara, Ahmed Imteaj and Abdur R. Shahid	1866
<i>RouteLMT: Learned Sample Routing for Hybrid LLM Translation Deployment</i>	
yingfeng luo, Hongyu Liu, DingYang Lin, Kaiyan Chang, Chenglong Wang, Bei Li, Quan Du, Tong Xiao and JingBo Zhu	1886
<i>Intent-Driven Semantic ID Generation for Grounded Conversational News Recommendation</i>	
Hongyang Su, Beibei Kong, Lei Cheng, Chengxiang Zhuo, Zang Li and Chenyun Yu	1898
<i>Development and Benchmarking of a Blended Human-AI Qualitative Research Assistant</i>	
Joseph Matveyenko, James Liu, John David Parsons, Ryan Brown, Alina I. Palimaru, Vipul Gupta and Prateek Puri	1917

<i>Entity Exchange in the Wild: A Diagnostic Study of LLM Based Real-World Conversational Entity Extraction</i>	Soumya Jain and Ayush Kumar	1933
<i>Agentic Context Strategies for Multi-Format Document Understanding: When Should Language Models Use Tools?</i>	Mansi Uniyal, Mukul Singh and Ryan Nadel	1956
<i>MM-BizRAG: Rethinking Multimodal Retrieval-Augmented Generation for General Purpose Enterprise Q&A</i>	Hanoz Bhatena, Parin Rajesh Jhaveri, Rohan Mittal, Prateek Singh, Aymen Kallala, Rachneet Kaur, Yiqiao Jin, Zhen Zeng, Adwait Ratnaparkhi and Denis Kochedykov	1967
<i>AesX: Enhance Your Images with Stunning Aesthetic Beauty</i>	Yuyan Chen, Zhendong Hou, Lei Xia, Jiahao Li, Zhuolin Ji and Zhixu Li	2001
<i>Transfer Learning for Generalizable Automated LLM Improvement Pipeline for IVR Navigation</i>	Vishal Sankar Ram, Jason Kushner, Manas Paldhe and Youngseo Son	2012
<i>Beyond Instruction Optimization: Multi-Agent Error-Driven Class Description Refinement for LLM-Based Classification</i>	Hamvir Dev, Shivam Ratnakant Mhaskar, Sasanka Vutla and Anup Pattnaik	2031
<i>TEN: Table Explicitization, Neurosymbolically</i>	Nikita Mehrotra, Aayush Kumar, Sumit Gulwani, Arjun Radhakrishna and Ashish Tiwari	2050
<i>RADAR: Risk-Aware Distilled Adaptive Routing for Efficient Short-Form Video Platform Ecosystem Governance</i>	Baoyu Jing, Zixuan Wang, Junwen Chen, Xin Dong and Bingfeng Deng	2087
<i>From Graph to Text and Back: Semantic Fidelity in Automated Industrial Knowledge Graphs</i>	Kamyar Zeinalipour, Silvia Severini, Alessia Borghini, Sara Cardarelli, Marco Maggini and Marco Gori	2095
<i>Learning Selective LLM Autonomy from Copilot Feedback in Enterprise Customer Support Workflows</i>	Nikita Borovkov, Elisei Rykov, Olga Tsymboi, Sergei Filimonov, Nikita Surnachev, Dmitry Bitman and Anatolii Potapov	2112
<i>Prompt-Level Distillation: A Non-Parametric Alternative to Model Fine-Tuning for Efficient Reasoning</i>	Sanket Badhe and Deep Shah	2131
<i>MTIVE: Multi-Task Image Verification Engine Using Vision-Language Models for E-commerce</i>	Yu-Tong Cao, Vishnu Prabhakaran, Arunita Das, Purav Aggarwal and Anoop Saladi	2148
<i>Measuring and Mitigating Racial Bias in Embedding Models: A Comparative Study for Law Enforcement Retrieval</i>	Archan Dutta	2160
<i>VishBox v2: A Multi-Agent System for Adaptive Voice Phishing Simulation</i>	Sungmi Park, Daon Choi, Yoonmo Yang, Hong Yunyi and Heedou Kim	2168
<i>ResoDiff-44k: High-Fidelity Cross-Lingual Speech and Singing Synthesis via Discrete Diffusion</i>	Gyanendra Das and Sai Satyam Jena	2183
<i>TingIS: Real-time Risk Event Discovery from Noisy Customer Incidents at Enterprise Scale</i>	Jun Wang, Ziyin Zhang, Rui Wang, Hang Yu, Peng Di and Rui Wang	2191

<i>LoCar: Localization-Aware Evaluation of In-Vehicle Assistants through Fine-Grained Sociolinguistic Control</i>	
Seogyeong Jeong, Kiwoong Park, Seyoung Song, Eunsu Kim, Ken E Friedl, Jaeho Kim and Alice Oh.....	2204
<i>Index-Time Prefix Injection for Multi-Tenant Retrieval: Improving Search Relevance Without Model Fine-Tuning</i>	
Vaibhav Varshney and Manjunatha Naik Mc	2231
<i>Benchmarking Testing in Automated Theorem Proving</i>	
Jongyoon Kim, Hojae Han and seung-won hwang	2241
<i>Graph-Based Phonetic Error Correction of Noisy ASR</i>	
Pratik Rakesh Singh, Mohammadi Zaki, Aneesh Mukkamla and Pankaj Wasnik	2261
<i>CAT: Confidence-Adaptive Thinking for Efficient Reasoning of Large Reasoning Models</i>	
Qizhi Jiang, Shuo Wang, Pei Ke, Yuhang Song and Ke Qin.....	2271
<i>LatentGate: Low-Latency Semantic Routing via Frozen-Backbone Probing of Small Language Models</i>	
Shivam Ratnakar, Abhiroop Talasila and Vinayak K Doifode	2284
<i>Thinking with Reasoning Skills: Fewer Tokens, More Accuracy</i>	
Guangxiang Zhao, Qilong Shi, Xusen Xiao, Xiangzheng Zhang, Tong Yang and Lin Sun...	2295