

Findings of EMNLP 2025

2025 Conference on Empirical Methods in
Natural Language Processing

Suzhou, China
4-9 November 2025

Volume 1 of 33

ISBN: 979-8-3313-2866-5

Printed from e-media with permission by:

Curran Associates, Inc.
57 Morehouse Lane
Red Hook, NY 12571



Some format issues inherent in the e-media version may also appear in this print version.

Copyright© (2025) by the Association for Computational Linguistics
All rights reserved.

Printed with permission by Curran Associates, Inc. (2026)

For permission requests, please contact the Association for Computational Linguistics
at the address below.

Association for Computational Linguistics
209 N. Eighth Street
Stroudsburg, Pennsylvania 18360

Phone: 1-570-476-8006
Fax: 1-570-476-0860

acl@aclweb.org

Additional copies of this publication are available from:

Curran Associates, Inc.
57 Morehouse Lane
Red Hook, NY 12571 USA
Phone: 845-758-0400
Fax: 845-758-2633
Email: curran@proceedings.com
Web: www.proceedings.com

Table of Contents

<i>Automating Alternative Generation in Decision-Making</i> Yevhen Kostyuk, Clara Seyfried and Chris Reed	1
<i>Bias Analysis and Mitigation through Protected Attribute Detection and Regard Classification</i> Takuma Udagawa, Yang Zhao, Hiroshi Kanayama and Bishwaranjan Bhattacharjee	16
<i>Large Language Models Might Not Care What You Are Saying: Prompt Format Beats Descriptions</i> Chenming Tang, Zhixiang Wang, Hao Sun and Yunfang Wu	26
<i>Boundary Matters: Leveraging Structured Text Plots for Long Text Outline Generation</i> Yuanchi Ma, Jiamou Liu, Hui He, Libo Zhang, Haoyuan Li and Zhendong Niu	49
<i>Can Large Language Models Personalize Dialogues to Generational Styles?</i> Pier Felice Balestrucci, Ondrej Dusek, Luca Anselma and Alessandro Mazzei	64
<i>Toward Optimal LLM Alignments Using Two-Player Games</i> Rui Zheng, Hongyi Guo, Zhihan Liu, Xiaoying Zhang, Yuanshun Yao, Xiaojun Xu, Zhaoran Wang, Zhiheng Xi, Tao Gui, Qi Zhang, Xuanjing Huang, Yang Liu and Hang Li	78
<i>Structural Patent Classification Using Label Hierarchy Optimization</i> Mengting Gui, Shufeng Hao, Chongyang Shi and Qi Zhang	100
<i>Exploring Hyperbolic Hierarchical Structure for Multimodal Rumor Detection</i> Md Mahbubur Rahman, Shufeng Hao, Chongyang Shi, An Lao and Jinyan Liu	115
<i>Multi-Surrogate-Objective Optimization for Neural Topic Models</i> Tue Le, Hoang Tran Vuong, Tung Nguyen, Linh Ngo Van, Dinh Viet Sang, Trung Le and Thien Huu Nguyen	135
<i>How Diversely Can Language Models Solve Problems? Exploring the Algorithmic Diversity of Model-Generated Code</i> Seonghyeon Lee, HeeJae Chon, Joonwon Jang, Dongha Lee and Hwanjo Yu	152
<i>ReAL: How Can LLMs Simulate the Real Teacher? Retrieval-enhanced Agent for Adaptive Learning</i> Rui Lv, Qi Liu, Weibo Gao, Jiatong Li, Kai Zhang and Shiwei Tong	168
<i>LLMsPark: A Benchmark for Evaluating Large Language Models in Strategic Gaming Contexts</i> Junhao Chen, Jingbo Sun, Xiang Li, Haidong Xin, Yuhao Xue, Yibin Xu and Hao Zhao	182
<i>Versatile Framework for Song Generation with Prompt-based Control</i> Yu Zhang, Wenxiang Guo, Changhao Pan, Zhiyuan Zhu, Ruiqi Li, Jingyu Lu, Rongjie Huang, Ruiyuan Zhang, Zhiqing Hong, Ziyue Jiang and Zhou Zhao	195
<i>InsBank: Evolving Instruction Subset for Ongoing Alignment</i> Jiayi Shi, Yiwei Li, Shaoxiong Feng, Peiwen Yuan, Xinglin Wang, Yueqi Zhang, Chuyi Tan, Boyuan Pan, Huan Ren, Yao Hu and Kan Li	220
<i>TL-Training: A Task-Feature-Based Framework for Training Large Language Models in Tool Use</i> Junjie Ye, Yilong Wu, Sixian Li, Yuming Yang, Zhiheng Xi, Tao Gui, Qi Zhang, Xuanjing Huang, Peng Wang, Zhongchao Shi, Jianping Fan and Zhengyin Du	239
<i>DCMKC: A Dual Consistency Matching Approach for Multi-hop Question Answering in LLMs</i> Xinyi Wang, Yiping Song, Chang Liu, Tingjin Luo, Bo Liu, Zheng Xie and Minlie Huang	259

<i>On Domain-Adaptive Post-Training for Multimodal Large Language Models</i>	
Daixuan Cheng, Shaohan Huang, Ziyu Zhu, Xintong Zhang, Xin Zhao, Zhongzhi Luan, Bo Dai and Zhenliang Zhang	274
<i>CPO: Addressing Reward Ambiguity in Role-playing Dialogue via Comparative Policy Optimization</i>	
Jing Ye, Rui Wang, Yuchuan Wu, Victor Ma, Feiteng Fang, Fei Huang and Yongbin Li	297
<i>SPPD: Self-training with Process Preference Learning Using Dynamic Value Margin</i>	
Hao Yi, Qingyang Li, Yulan Hu, Fuzheng Zhang, Di Zhang and Yong Liu	324
<i>Error Classification of Large Language Models on Math Word Problems: A Dynamically Adaptive Framework</i>	
Zhangyue Yin, YuHong Sun, Xuanjing Huang, Xipeng Qiu and Hui Zhao	338
<i>sudoLLM: On Multi-role Alignment of Language Models</i>	
Soumadeep Saha, Akshay Chaturvedi, Joy Mahapatra and Utpal Garain	366
<i>DAC: Decomposed Automation Correction for Text-to-SQL</i>	
Dingzirui Wang, Longxu Dou, Xuanliang Zhang, Qingfu Zhu and Wanxiang Che	385
<i>VehicleWorld: A Highly Integrated Multi-Device Environment for Intelligent Vehicle Interaction</i>	
Jie Yang, Jiajun Chen, Zhangyue Yin, Shuo Chen, Yuxin Wang, Yiran Guo, Yuan Li, Yining Zheng, Xuanjing Huang and Xipeng Qiu	403
<i>End-to-End Optimization for Multimodal Retrieval-Augmented Generation via Reward Backpropagation</i>	
Zhiyuan Fan, Longfei Yun, Ming Yan, Yumeng Wang, Dadi Guo, Brian Mak, James Kwok and Yi R. Fung	443
<i>Audio-Aware Large Language Models as Judges for Speaking Styles</i>	
Cheng-Han Chiang, Xiaofei Wang, Chung-Ching Lin, Kevin Lin, Linjie Li, Radu Kopetz, Yao Qian, Zhendong Wang, Zhengyuan Yang, Hung-yi Lee and Lijuan Wang	467
<i>Evaluation of Text-to-Image Generation from a Creativity Perspective</i>	
Xinhao Wang, Xinyu Ma, ShengYong Ding and Derek F. Wong	481
<i>Perovskite-LLM: Knowledge-Enhanced Large Language Models for Perovskite Solar Cell Research</i>	
Xiang Liu, Penglei Sun, Shuyan Chen, Longhan Zhang, Peijie Dong, Huajie You, Yongqi Zhang, Chang Yan, Xiaowen Chu and Tong-yi Zhang	494
<i>ProPy: Building Interactive Prompt Pyramids upon CLIP for Partially Relevant Video Retrieval</i>	
Yi Pan, Yujia Zhang, Michael Kampffmeyer and Xiaoguang Zhao	519
<i>Multilingual Datasets for Custom Input Extraction and Explanation Requests Parsing in Conversational XAI Systems</i>	
Qianli Wang, Tatiana Anikina, Nils Feldhus, Simon Ostermann, Fedor Splitt, Jiaao Li, Yoana Tsoneva, Sebastian Möller and Vera Schmitt	534
<i>Toolscaler: Scalable Generative Tool Calling via Structure-Aware Semantic Tokenization</i>	
Yunyue Su, Zhang Jinshuai, Bowen Fang, Wen Ye, Jinghao Zhang, Bowen Song, Weiqiang Wang, Qiang Liu and Liang Wang	556
<i>LaMP-Val: Large Language Models Empower Personalized Valuation in Auction</i>	
Jie Sun, Tianyu Zhang, Houcheng Jiang, Kexin Huang, Xiang Shu, Zhibo Zhu, Lintao Ma, Xingyu Lu, Jun Zhou, Junkang Wu, Chi Luo, An Zhang, Jiancan Wu and Xiang Wang	579

<i>Exploring Model Kinship for Merging Large Language Models</i> Yedi Hu, Yunzhi Yao, Ningyu Zhang, Huajun Chen and Shumin Deng	596
<i>MULTITAT: Benchmarking Multilingual Table-and-Text Question Answering</i> Xuanliang Zhang, Dingzirui Wang, Keyan Xu, Qingfu Zhu and Wanxiang Che	626
<i>LoRA-MGPO: Mitigating Double Descent in Low-Rank Adaptation via Momentum-Guided Perturbation Optimization</i> Yupeng Chang, Chenlu Guo, Yi Chang and Yuan Wu	648
<i>R-LoRA: Randomized Multi-Head LoRA for Efficient Multi-task Learning</i> Jinda Liu, Yi Chang and Yuan Wu	660
<i>RACQC: Advanced Retrieval-Augmented Generation for Chinese Query Correction</i> Jinbo Su, Lingzhe Gao, Wei Li, Shihao Liu, Haojie Lei, Xinyi Wang, Yuanzhao Guo, Ke Wang, Daiting Shi and Dawei Yin	675
<i>Mechanistic Understanding and Mitigation of Language Confusion in English-Centric Large Language Models</i> Ercong Nie, Helmut Schmid and Hinrich Schuetze	690
<i>Assessing and Mitigating Medical Knowledge Drift and Conflicts in Large Language Models</i> Weiyi Wu, Xinwen Xu, Chongyang Gao, Xingjian Diao, Siting Li, Lucas A. Salas and Jiang Gui	707
<i>Improving LLM Reasoning through Interpretable Role-Playing Steering</i> Anyi Wang, Dong Shu, Yifan Wang, Yunpu Ma and Mengnan Du	731
<i>R2A-TLS: Reflective Retrieval-Augmented Timeline Summarization with Causal-Semantic Integration</i> Chenlong Bao, Shijie Li, Minghao Hu, Ming Qiao, Bin Zhang, Jin-Tao Tang, Shasha Li and Ting Wang	752
<i>MedEBench: Diagnosing Reliability in Text-Guided Medical Image Editing</i> Minghao Liu, Zhitao He, Zhiyuan Fan, Qingyun Wang and Yi R. Fung	767
<i>FairCoT: Enhancing Fairness in Text-to-Image Generation via Chain of Thought Reasoning with Multimodal Large Language Models</i> Zahraa Al Sahili, Ioannis Patras and Matthew Purver	792
<i>Bag of Tricks for Sparse Mixture-of-Experts: A Benchmark Across Reasoning, Efficiency, and Safety</i> Mufan Qiu, Zheyu Shen, Pingzhi Li, Ang Li and Tianlong Chen	817
<i>Don't Take the Premise for Granted: Evaluating the Premise Critique Ability of Large Language Models</i> Jinzhe Li, Gengxu Li, Yi Chang and Yuan Wu	836
<i>Mitigating Geospatial Knowledge Hallucination in Large Language Models: Benchmarking and Dynamic Factuality Aligning</i> Shengyuan Wang, Jie Feng, Tianhui Liu, Dan Pei and Yong Li	870
<i>The Power of Framing: How News Headlines Guide Search Behavior</i> Amrit Poudel, Maria Milkowski and Tim Wenginger	889
<i>DivLogicEval: A Framework for Benchmarking Logical Reasoning Evaluation in Large Language Models</i> Tsz Ting Chung, Lemao Liu, Mo Yu and Dit-Yan Yeung	901

<i>THCM-CAL: Temporal-Hierarchical Causal Modelling with Conformal Calibration for Clinical Risk Prediction</i>	
Xin Zhang, Qiyu Wei, Yingjie Zhu, Fanyi Wu and Sophia Ananiadou	916
<i>GenPilot: A Multi-Agent System for Test-Time Prompt Optimization in Image Generation</i>	
Wen Ye, Zhaocheng Liu, Gui Yuwei, Tingyu Yuan, Yunyue Su, Bowen Fang, Chaoyang Zhao, Qiang Liu and Liang Wang	929
<i>Grounded-VideoLLM: Sharpening Fine-grained Temporal Grounding in Video Large Language Models</i>	
Haibo Wang, Zhiyang Xu, Yu Cheng, Shizhe Diao, Yufan Zhou, Yixin Cao, Qifan Wang, Weifeng Ge and Lifu Huang	959
<i>DongbaMIE: A Multimodal Information Extraction Dataset for Evaluating Semantic Understanding of Dongba Pictograms</i>	
Xiaojun Bi, Shuo Li, Junyao Xing, Ziyue Wang, Fuwen Luo, Weizheng Qiao, Lu Han, Ziwei Sun, Peng Li and Yang Liu	976
<i>Optimizing Cross-Client Domain Coverage for Federated Instruction Tuning of Large Language Models</i>	
Zezhou Wang, Yaxin Du, Xingjun Ma, Yu-Gang Jiang, Zhuzhong Qian and Siheng Chen . . .	991
<i>Aligning Black-Box LLMs for Aspect Sentiment Quad Prediction</i>	
Shichen Li, Jiawei Zhang, Zhongqing Wang and Peifeng Li	1012
<i>Multifaceted Evaluation of Audio-Visual Capability for MLLMs: Effectiveness, Efficiency, Generalizability and Robustness</i>	
Yusheng Zhao, Xiao Luo, Junyu Luo, Weizhi Zhang, Zhiping Xiao, Wei Ju, Philip S. Yu and Ming Zhang	1026
<i>Two Steps from Hell: Compositionality on Chemical LMs</i>	
Veronika Ganeeva, Kuzma Khrabrov, Artur Kadurin and Elena Tutubalina	1042
<i>GTA: Supervised-Guided Reinforcement Learning for Text Classification with Large Language Models</i>	
Min Zeng, Jingfei Sun, Xueyou Luo, Shiqi Zhang, Li Xie, Caiquan Liu and Xiaoxin Chen . .	1050
<i>Unearthing Gems from Stones: Policy Optimization with Negative Sample Augmentation for LLM Reasoning</i>	
Zhaohui Yang, Yuxiao Ye, Shilei Jiang, Shihong Deng, Chen Hu, Linjing Li and Daxin Jiang	1061
<i>LEAF: Large Language Diffusion Model for Time Series Forecasting</i>	
Yuhang Pei, Tao Ren, Yifan Wang, Zhipeng Sun, Wei Ju, Chong Chen, Xian-Sheng Hua and Xiao Luo	1076
<i>SPFT-SQL: Enhancing Large Language Model for Text-to-SQL Parsing by Self-Play Fine-Tuning</i>	
Yuhao Zhang, Shaoming Duan, Jinhang Su, Chuanyi Liu and Peiyi Han	1092
<i>Multilingual Verbalisation of Knowledge Graphs</i>	
Yifei Song, William Soto Martinez, Anna Nikiforovskaya, Evan Parker Kelly Chapple and Claire Gardent	1111
<i>LAGCL4Rec: When LLMs Activate Interactions Potential in Graph Contrastive Learning for Recommendation</i>	
Leqi Zheng, Chaokun Wang, Canzhi Chen, Jiajun Zhang, Cheng Wu, Zixin Song, Shannan Yan, Ziyang Liu and Hongwei Li	1163
<i>English as Defense Proxy: Mitigating Multilingual Jailbreak via Eliciting English Safety Knowledge</i>	
Zekai Zhang, Yiduo Guo, Jiuheng Lin, Shanghaoran Quan, Huishuai Zhang and Dongyan Zhao	1185

<i>Dagger Behind Smile: Fool LLMs with a Happy Ending Story</i>	
Xurui Song, Zhixin Xie, Shuo Huai, Jiayi Kong and Jun Luo	1197
<i>Mitigating Object Hallucinations in MLLMs via Multi-Frequency Perturbations</i>	
Shuo Li, Jiajun Sun, Guodong Zheng, Xiaoran Fan, Yujiong Shen, Yi Lu, Zhiheng Xi, Yuming Yang, Wenming Tan, Tao Ji, Tao Gui, Qi Zhang and Xuanjing Huang	1230
<i>Natural Context Drift Undermines the Natural Language Understanding of Large Language Models</i>	
Yulong Wu, Viktor Schlegel and Riza Batista-Navarro	1248
<i>Minimal Ranks, Maximum Confidence: Parameter-efficient Uncertainty Quantification for LoRA</i>	
Patryk Marszałek, Klaudia Bałazy, Jacek Tabor and Tomasz Kuśmierczyk	1260
<i>Chain-of-Thought Prompting Obscures Hallucination Cues in Large Language Models: An Empirical Evaluation</i>	
Jiahao Cheng, Tiancheng Su, Jia Yuan, Guoxiu He, Jiawei Liu, Xinqi Tao, Jingwen Xie and Huaxia Li	1272
<i>Large Language Model Evaluation via Matrix Nuclear-Norm</i>	
Yahan Li, Tingyu Xia, Yuan Wu and Yi Chang	1306
<i>From Grounding to Manipulation: Case Studies of Foundation Model Integration in Embodied Robotic Systems</i>	
Xiuchao Sui, Daiying Tian, Qi Sun, Ruirui Chen, Dongkyu Choi, Kenneth Kwok and Soujanya Poria	1324
<i>Flexible Thinking for Multimodal Emotional Support Conversation via Reinforcement Learning</i>	
Fanfan Wang, Xiangqing Shen, Jianfei Yu and Rui Xia	1341
<i>ORAL: Prompting Your Large-Scale LoRAs via Conditional Recurrent Diffusion</i>	
Rana Shahroz, Dongwen Tang, Pingzhi Li, Kai Wang and Tianlong Chen	1357
<i>NLoRA: Nyström-Initiated Low-Rank Adaptation for Large Language Models</i>	
Chenlu Guo, Yi Chang and Yuan Wu	1371
<i>Bhaasha, Bhāṣā, Zaban: A Survey for Low-Resourced Languages in South Asia – Current Stage and Challenges</i>	
Sampoorna Poria and Xiaolei Huang	1386
<i>DISCO Balances the Scales: Adaptive Domain- and Difficulty-Aware Reinforcement Learning on Imbalanced Data</i>	
Yuhang Zhou, Jing Zhu, Shengyi Qian, Zhuokai Zhao, Xiyao Wang, Xiaoyu Liu, Ming Li, Paiheng Xu, Wei Ai and Furong Huang	1407
<i>What Makes for Good Image Captions?</i>	
Delong Chen, Samuel Cahyawijaya, Etsuko Ishii, Ho Shu Chan, Yejin Bang and Pascale Fung	1420
<i>What’s Not Said Still Hurts: A Description-Based Evaluation Framework for Measuring Social Bias in LLMs</i>	
Jinhao Pan, Chahat Raj, Ziyu Yao and Ziwei Zhu	1438
<i>Identifying Rare Languages in Common Crawl Data is a Needles-in-a-Haystack Problem</i>	
Rasul Dent, Pedro Ortiz Suarez, Thibault Clérice and Benoît Sagot	1460
<i>Training Language Models to Critique With Multi-agent Feedback</i>	
Tian Lan, Wenwei Zhang, Chengqi Lyu, Shuaibin Li, Chen Xu, Heyan Huang, Dahua Lin, Xian-Ling Mao and Kai Chen	1474

<i>RELIC: Enhancing Reward Model Generalization for Low-Resource Indic Languages with Few-Shot Examples</i>	
Soumya Suvra Ghosal, Vaibhav Singh, Akash Ghosh, Soumyabrata Pal, Subhadip Baidya, Sriparna Saha and Dinesh Manocha	1502
<i>Invoke Interfaces Only When Needed: Adaptive Invocation for Large Language Models in Question Answering</i>	
Jihao Zhao, Chunlai Zhou, Daixuan Li, Shuaishuai Zu and Biao Qin	1518
<i>SQLSpace: A Representation Space for Text-to-SQL to Discover and Mitigate Robustness Gaps</i>	
Neha Srikanth, Victor Bursztyn, Puneet Mathur and Ani Nenkova	1533
<i>One More Modality: Does Abstract Meaning Representation Benefit Visual Question Answering?</i>	
Abhidip Bhattacharyya, Emma Markle and Shira Wein	1560
<i>DP-GTR: Differentially Private Prompt Protection via Group Text Rewriting</i>	
Mingchen Li, Heng Fan, Song Fu, Junhua Ding and Yunhe Feng	1573
<i>Legal Mathematical Reasoning with LLMs: Procedural Alignment through Two-Stage Reinforcement Learning</i>	
Kepu Zhang, Guofu Xie, Weijie Yu, Mingyue Xu, Xu Tang, Yaxin Li and Jun Xu	1586
<i>ModelingAgent: Bridging LLMs and Mathematical Modeling for Real-World Challenges</i>	
Cheng Qian, Hongyi Du, Hongru Wang, Xiushi Chen, Yuji Zhang, Avirup Sil, ChengXiang Zhai, Kathleen McKeown and Heng Ji	1599
<i>Beyond Coarse Labels: Fine-Grained Problem Augmentation and Multi-Dimensional Feedback for Emotional Support Conversation</i>	
Yuanchen Shi, Jiawang Hao and Fang Kong	1634
<i>FinHEAR: Human Expertise and Adaptive Risk-Aware Temporal Reasoning for Financial Decision-Making</i>	
Jiaxiang Chen, Mingxi Zou, Zhuo Wang, Qifan Wang, Danny Dongning Sun, Zhang Chi and Zenglin Xu	1648
<i>EvoKV: Evolutionary KV Cache Compression for LLM Inference</i>	
Bohan Yu and Yekun Chai	1673
<i>A Survey on Sparse Autoencoders: Interpreting the Internal Mechanisms of Large Language Models</i>	
Dong Shu, Xuansheng Wu, Haiyan Zhao, Daking Rai, Ziyu Yao, Ninghao Liu and Mengnan Du	1690
<i>Large Vision-Language Model Alignment and Misalignment: A Survey Through the Lens of Explainability</i>	
Dong Shu, Haiyan Zhao, Jingyu Hu, Weiru Liu, Ali Payani, Lu Cheng and Mengnan Du . . .	1713
<i>Attention Consistency for LLMs Explanation</i>	
Tian Lan, Jinyuan XU, Xue HE, Jenq-Neng Hwang and Lei Li	1736
<i>Confusion is the Final Barrier: Rethinking Jailbreak Evaluation and Investigating the Real Misuse Threat of LLMs</i>	
Yu Yan, Sheng Sun, Zhe Wang, Yijun Lin, Zenghao Duan, Zhifei Zheng, Min Liu, Zhiyi Yin and Jianping Zhang	1751
<i>CCL-XCoT: An Efficient Cross-Lingual Knowledge Transfer Method for Mitigating Hallucination Generation</i>	
Zheng Weihua, Roy Ka-Wei Lee, Zhengyuan Liu, Wu Kui, AiTi Aw and Bowei Zou	1768

<i>Evaluating Step-by-step Reasoning Traces: A Survey</i> Jinu Lee and Julia Hockenmaier	1789
<i>Beyond Guilt: Legal Judgment Prediction with Trichotomous Reasoning</i> Kepu Zhang, Haoyue Yang, Xu Tang, Weijie Yu and Jun Xu	1815
<i>Not Every Token Needs Forgetting: Selective Unlearning Balancing Forgetting and Utility in Large Language Models</i> Yixin Wan, Anil Ramakrishna, Kai-Wei Chang, Volkan Cevher and Rahul Gupta	1827
<i>DisastIR: A Comprehensive Information Retrieval Benchmark for Disaster Management</i> Kai Yin, Xiangjue Dong, Chengkai Liu, Lipai Huang, Yiming Xiao, Zhewei Liu, Ali Mostafavi and James Caverlee	1836
<i>Data or Language Supervision: What Makes CLIP Better than DINO?</i> Yiming Liu, Yuhui Zhang, Dhruva Ghosh, Ludwig Schmidt and Serena Yeung-Levy	1868
<i>Do LLMs Understand Wine Descriptors Across Cultures? A Benchmark for Cultural Adaptations of Wine Reviews</i> Chenye Zou, Xingyue Wen, Tianyi Hu, Qian Janice Wang and Daniel Hershcovich	1875
<i>DeFT-X: Denoised Sparse Fine-Tuning for Zero-Shot Cross-Lingual Transfer</i> Sona Elza Simon and Preethi Jyothi	1895
<i>Memory-enhanced Large Language Model for Cross-lingual Dependency Parsing via Deep Hierarchical Syntax Understanding</i> Jianjian Liu, Ying Li, Zhengtao Yu, Shun Su, Shengxiang Gao and Yuxin Huang	1910
<i>Developing and Utilizing a Large-Scale Cantonese Dataset for Multi-Tasking in Large Language Models</i> Jiyue Jiang, Alfred Kar Yin Truong, Yanyu Chen, Qinghang Bao, Sheng Wang, Pengan Chen, Jiuming Wang, Lingpeng Kong, Yu Li and Chuan Wu	1924
<i>A Structured Framework for Evaluating and Enhancing Interpretive Capabilities of Multimodal LLMs in Culturally Situated Tasks</i> Haorui Yu, Ramon Ruiz-Dolz and Qiufeng Yi	1945
<i>Train a Unified Multimodal Data Quality Classifier with Synthetic Data</i> Weizhi Wang, Rongmei Lin, Shiyang Li, Colin Lockard, Ritesh Sarkhel, Sanket Lokegaonkar, Jingbo Shang, Xifeng Yan, Nasser Zalmout and Xian Li	1972
<i>Self-Improvement in Multimodal Large Language Models: A Survey</i> Shijian Deng, Kai Wang, Tianyu Yang, Harsh Singh and Yapeng Tian	1987
<i>Towards Achieving Concept Completeness for Textual Concept Bottleneck Models</i> Milan Bhan, Yann Choho, Jean-Noël Vittaut, Nicolas Chesneau, Pierre Moreau and Marie-Jeanne Lesot	2007
<i>EmoBench-UA: A Benchmark Dataset for Emotion Detection in Ukrainian</i> Daryna Dementieva, Nikolay Babakov and Alexander Fraser	2025
<i>Scientific Paper Retrieval with LLM-Guided Semantic-Based Ranking</i> Yunyi Zhang, Ruozhen Yang, Siqi Jiao, SeongKu Kang and Jiawei Han	2049
<i>DLIR: Spherical Adaptation for Cross-Lingual Knowledge Transfer of Sociological Concepts Alignment</i> Zeqiang Wang, Jon Johnson and Suparna De	2061

<i>Test-Time Steering for Lossless Text Compression via Weighted Product of Experts</i>	
Qihang Zhang, Muchen Li, Ziao Wang, Renjie Liao and Lele Wang	2076
<i>Zero-Shot Contextual Embeddings via Offline Synthetic Corpus Generation</i>	
Philip Lippmann and Jie Yang	2089
<i>The Hallucination Tax of Reinforcement Finetuning</i>	
Linxin Song, Taiwei Shi and Jieyu Zhao	2105
<i>Tracing Multilingual Factual Knowledge Acquisition in Pretraining</i>	
Yihong Liu, Mingyang Wang, Amir Hossein Kargaran, Felicia Körner, Ercong Nie, Barbara Plank, François Yvon and Hinrich Schuetze	2121
<i>Exploring the Vulnerability of the Content Moderation Guardrail in Large Language Models via Intent Manipulation</i>	
Jun Zhuang, Haibo Jin, Ye Zhang, Zhengjian Kang, Wenbin Zhang, Gaby G. Dagher and Haohan Wang	2147
<i>Examining Multilingual Embedding Models Cross-Lingually Through LLM-Generated Adversarial Examples</i>	
Andrianos Michail, Simon Clematide and Rico Sennrich	2161
<i>EmoGist: Efficient In-Context Learning for Visual Emotion Understanding</i>	
Ronald Seoh and Dan Goldwasser	2171
<i>Soft Token Attacks Cannot Reliably Audit Unlearning in Large Language Models</i>	
Haokun Chen, Sebastian Szyller, Weilin Xu and Nageen Himayat	2183
<i>Bridging the Editing Gap in LLMs: FineEdit for Precise and Targeted Text Modifications</i>	
Yiming Zeng, Wanhao Yu, Zexin Li, Tao Ren, Yu Ma, Jinghan Cao, Xiyan Chen and Tingting Yu	2193
<i>LLM-based Conversational Recommendation Agents with Collaborative Verbalized Experience</i>	
Yaochen Zhu, Harald Steck, Dawen Liang, Yinhan He, Nathan Kallus and Jundong Li	2207
<i>Hardware-Aware Parallel Prompt Decoding for Memory-Efficient Acceleration of LLM Inference</i>	
Hao Mark Chen, Wayne Luk, Yiu Ka Fai Cedric, Rui Li, Konstantin Mishchenko, Stylianos Venieris and Hongxiang Fan	2221
<i>Measuring Sycophancy of Language Models in Multi-turn Dialogues</i>	
Jiseung Hong, Grace Byun, Seungone Kim and Kai Shu	2239
<i>On the Role of Entity and Event Level Conceptualization in Generalizable Reasoning: A Survey of Tasks, Methods, Applications, and Future Directions</i>	
Weiqi Wang, Tianqing Fang, Haochen Shi, Baixuan Xu, Wenxuan Ding, Liyu Zhang, Wei Fan, Jiabin Bai, Haoran Li, Xin Liu and Yangqiu Song	2260
<i>Mitigating Visual Knowledge Forgetting in MLLM Instruction-tuning via Modality-decoupled Gradient Descent</i>	
Junda Wu, Yuxin Xiong, Xintong Li, Yu Xia, Ruoyu Wang, Yu Wang, Tong Yu, Sungchul Kim, Ryan A. Rossi, Lina Yao, Jingbo Shang and Julian McAuley	2282
<i>PathoHR: Hierarchical Reasoning for Vision-Language Models in Pathology</i>	
Yating Huang, Ziyang Huang, Lintao Xiang, Qijun Yang and Hujun Yin	2296

<i>What's Up, Doc?": Analyzing How Users Seek Health Information in Large-Scale Conversational AI Datasets</i>	
Akshay Paruchuri, Maryam Aziz, Rohit Vartak, Ayman Ali, Best Uchehara, Xin Liu, Ishan Chatterjee and Monica Agrawal	2312
<i>Dynamic Evaluation for Oversensitivity in LLMs</i>	
Sophia Xiao Pu, Sitao Cheng, Xin Eric Wang and William Yang Wang	2337
<i>Self-Correcting Code Generation Using Small Language Models</i>	
Jeonghun Cho, Deokhyung Kang, Hyounghun Kim and Gary Lee	2345
<i>A Unified Framework for N-ary Property Information Extraction in Materials Science</i>	
Van-Thuy Phi and Yuji Matsumoto	2369
<i>A Benchmark for Translations Across Styles and Language Variants</i>	
Xin Tan, Bowei Zou and AiTi Aw	2389
<i>ManuSearch: Democratizing Deep Search in Large Language Models with a Transparent and Open Multi-Agent Framework</i>	
Lisheng Huang, Yichen Liu, Jinhao Jiang, Rongxiang Zhang, Jiahao Yan, Junyi Li and Xin Zhao	2403
<i>Proactive User Information Acquisition via Chats on User-Favored Topics</i>	
Shiki Sato, Jun Baba, Asahi Hentona, Shinji Iwata, Akifumi Yoshimoto and Koichiro Yoshino	2418
<i>Evaluating Text Generation Quality Using Spectral Distances of Surprisal</i>	
Zhichen Liu, Yongyuan Li, Yang Xu, Yu Wang, Yingfang Yuan and Zuhao Yang	2444
<i>NLP-ADBench: NLP Anomaly Detection Benchmark</i>	
Yuangang Li, Jiaqi li, Zhuo Xiao, Tiankai Yang, Yi Nian, Xiyang Hu and Yue Zhao	2464
<i>Toward Inclusive Language Models: Sparsity-Driven Calibration for Systematic and Interpretable Mitigation of Social Biases in LLMs</i>	
Prommy Sultana Hossain, Chahat Raj, Ziwei Zhu, Jessica Lin and Emanuela Marasco	2475
<i>Table-Text Alignment: Explaining Claim Verification Against Tables in Scientific Papers</i>	
Xanh Ho, Sunisth Kumar, Yun-Ang Wu, Florian Boudin, Atsuhiko Takasu and Akiko Aizawa	2509
<i>DCRM: A Heuristic to Measure Response Pair Quality in Preference Optimization</i>	
Chengyu Huang and Tanya Goyal	2518
<i>Advancing Reasoning with Off-the-Shelf LLMs: A Semantic Structure Perspective</i>	
Pengfei He, Zitao Li, Yue Xing, Yaliang Li, Jiliang Tang and Bolin Ding	2538
<i>LLM-based Open Domain Planning by Leveraging Entity-Attribute-Level Domain Models</i>	
Dongning Rao, Songlin He, Zhihua Jiang and Ruishi Liang	2567
<i>DICP: Deep In-Context Prompt for Event Causality Identification</i>	
Lin Mu, Jun Shen, Li Ni, Lei Sang, Zhize Wu, Peiquan Jin and Yiwen Zhang	2589
<i>Seeing is Believing: Emotion-Aware Audio-Visual Language Modeling for Expressive Speech Generation</i>	
Weiting Tan, Jiachen Lian, Hirofumi Inaguma, Paden Tomasello, Philipp Koehn and Xutai Ma	2600

<i>GRV-KBQA: A Three-Stage Framework for Knowledge Base Question Answering with Decoupled Logical Structure, Semantic Grounding and Structure-Aware Validation</i>	
Yuhang Tian, Pan Yang, Dandan Song, Zhijing Wu and Hao Wang	2618
<i>Improving Prompt Generalization for Cross-prompt Essay Trait Scoring from the Scoring-invariance Perspective</i>	
Jiong Wang and Shengquan Yu	2633
<i>When Format Changes Meaning: Investigating Semantic Inconsistency of Large Language Models</i>	
Cheongwoong Kang, Jongeun Baek, Yeonjea Kim and Jaesik Choi	2647
<i>ASTPrompter: Preference-Aligned Automated Language Model Red-Teaming to Generate Low-Perplexity Unsafe Prompts</i>	
Amelia Hardy, Houjun Liu, Allie Griffith, Bernard Lange, Duncan Eddy and Mykel Kochenderfer	2668
<i>How Do Large Language Models Perform on PDE Discovery: A Coarse-to-fine Perspective</i>	
Xiao Luo, Changhu Wang, Yizhou Sun and Wei Wang	2684
<i>Rethinking Data Selection at Scale: Random Selection is Almost All You Need</i>	
Tingyu Xia, Bowen Yu, Kai Dang, An Yang, Yuan Wu, Yuan Tian, Yi Chang and Junyang Lin	2698
<i>PromptKeeper: Safeguarding System Prompts for LLMs</i>	
Zhifeng Jiang, Zhihua Jin and Guoliang He	2712
<i>Automating eHMI Action Design with LLMs for Automated Vehicle Communication</i>	
Ding Xia, Xinyue Gui, Fan Gao, Dongyuan Li, Mark Colley and Takeo Igarashi	2729
<i>A Dynamic Fusion Model for Consistent Crisis Response</i>	
Xiaoying Song, Anirban Saha Anik, Eduardo Blanco, Vanessa Frias-Martinez and Lingzi Hong	2753
<i>UIOrchestra: Generating High-Fidelity Code from UI Designs with a Multi-agent System</i>	
Chuhuai Yue, Jiajun Chai, Yufei Zhang, Zixiang Ding, Xihao Liang, Peixin Wang, Shihai Chen, Wang Yixuan, Wangyanping, Guojun Yin and Wei Lin	2769
<i>CrossQG: Improving Difficulty-Controllable Question Generation through Consistency Enhancement</i>	
Kunze Li and Yu Zhang	2783
<i>Progressive Facial Granularity Aggregation with Bilateral Attribute-based Enhancement for Face-to-Speech Synthesis</i>	
Yejin Jeon, Youngjae Kim, Jihyun Lee, Hyoungun Kim and Gary Lee	2799
<i>Speaking at the Right Level: Literacy-Controlled Counterspeech Generation with RAG-RL</i>	
Xiaoying Song, Anirban Saha Anik, Dibakar Barua, Pengcheng Luo, Junhua Ding and Lingzi Hong	2812
<i>FNSCC: Fuzzy Neighborhood-Aware Self-Supervised Contrastive Clustering for Short Text</i>	
Zijian Zheng, Yonghe Lu and Jian Yin	2831
<i>AuraDial: A Large-Scale Human-Centric Dialogue Dataset for Chinese AI Psychological Counseling</i>	
Xiantao Zhang	2847
<i>TS-SQL: Test-driven Self-refinement for Text-to-SQL</i>	
Wenbo Xu, Haifeng Zhu, Liang Yan, Chuanyi Liu, Peiyi Han, Shaoming Duan and Jeff Z. Pan	2864

<i>DemonAgent: Dynamically Encrypted Multi-Backdoor Implantation Attack on LLM-based Agent</i> Pengyu Zhu, Zhenhong Zhou, Yuanhe Zhang, Shilinlu Yan, Kun Wang and Sen Su	2890
<i>MotivGraph-SoIQ: Integrating Motivational Knowledge Graphs and Socratic Dialogue for Enhanced LLM Ideation</i> Xinping Lei, Tong Zhou, Yubo Chen, Kang Liu and Jun Zhao	2913
<i>ExpertGenQA: Open-ended QA generation in Specialized Domains</i> Haz Sameen Shahgir, Chansong Lim, Jia Chen, Evangelos E. Papalexakis and Yue Dong	2934
<i>VisCoder: Fine-Tuning LLMs for Executable Python Visualization Code Generation</i> Yuansheng Ni, Ping Nie, Kai Zou, Xiang Yue and Wenhua Chen	2956
<i>Conversational Education at Scale: A Multi-LLM Agent Workflow for Procedural Learning and Pedagogic Quality Assessment</i> Jiahuan Pei, Fanghua Ye, Xin Sun, Wentao Deng, Koen Hindriks and Junxiao Wang	2984
<i>Visual Program Distillation with Template-Based Augmentation</i> Michal Shlapentokh-Rothman, Yu-Xiong Wang and Derek Hoiem	2998
<i>NeighXLM: Enhancing Cross-Lingual Transfer in Low-Resource Languages via Neighbor-Augmented Contrastive Pretraining</i> Sicheng Wang, Wenyi Wu and Zibo Zhang	3019
<i>ICLER: Intent Classification with Enhanced Reasoning</i> Dezheng Gao, Dong Xiaozheng, SHuangtao Yang and Bo Fu	3031
<i>PreGenie: An Agentic Framework for High-quality Visual Presentation Generation</i> Xiaojie Xu, Xinli Xu, Sirui Chen, Haoyu Chen, Fan Zhang and Ying-Cong Chen	3045
<i>RIVAL: Reinforcement Learning with Iterative and Adversarial Optimization for Machine Translation</i> Tianjiao Li, Mengran Yu, Chenyu Shi, Yanjun Zhao, Xiaojing Liu, Qi Zhang, Xuanjing Huang, Qiang Zhang and Jiayin Wang	3064
<i>MRAG: A Modular Retrieval Framework for Time-Sensitive Question Answering</i> Siye Zhang, Yuxiang Xue, Yiming Zhang, Xiaobao Wu, Anh Tuan Luu and Chen Zhao	3080
<i>CoT-RAG: Integrating Chain of Thought and Retrieval-Augmented Generation to Enhance Reasoning in Large Language Models</i> Feiyang Li, Peng Fang, Zhan Shi, Arijit Khan, Fang Wang, Weihao Wang, Zhangxin-hw and Cui Yongjian	3119
<i>TabDSR: Decompose, Sanitize, and Reason for Complex Numerical Reasoning in Tabular Data</i> Changjiang Jiang, Fengchang Yu, Haihua Chen, Wei Lu and Jin Zeng	3172
<i>Chain-of-Thought Matters: Improving Long-Context Language Models with Reasoning Path Supervision</i> Dawei Zhu, Xiyu Wei, Guangxiang Zhao, Wenhao Wu, Haosheng Zou, Junfeng Ran, XWang, Lin Sun, Xiangzheng Zhang and Sujian Li	3197
<i>Multimodal Document-level Triple Extraction via Dynamic Graph Enhancement and Relation-Aware Reflection</i> Xiang Li, Runhai Jiao, Zhou Changyu, Shoupeng Qiao, Ruojiao Qiao and Ruifan Li	3212
<i>Distill Visual Chart Reasoning Ability from LLMs to MLLMs</i> Wei He, Zhiheng Xi, Wanxu Zhao, Xiaoran Fan, Yiwen Ding, Zifei Shan, Tao Gui, Qi Zhang and Xuanjing Huang	3224

<i>FlowMalTrans: Unsupervised Binary Code Translation for Malware Detection Using Flow-Adapter Architecture</i>	
Minghao Hu, Junzhe Wang, Weisen Zhao, Qiang Zeng and Lannan Luo	3251
<i>AdaTP: Attention-Debiased Token Pruning for Video Large Language Models</i>	
Fengyuan Sun, Leqi Shen, Hui Chen, Sicheng Zhao, Jungong Han and Guiguang Ding	3273
<i>AdaptFlow: Adaptive Workflow Optimization via Meta-Learning</i>	
Runchuan Zhu, Bowen Jiang, Lingrui Mei, Fangkai Yang, Lu Wang, Haoxiang Gao, Fengshuo Bai, Pu Zhao, Qingwei Lin, Saravan Rajmohan and Dongmei Zhang	3287
<i>LMUNIT: Fine-grained Evaluation with Natural Language Unit Tests</i>	
Jon Saad-Falcon, Rajan Pathe Vivek, William Berrios, Nandita Shankar Naik, Matija Franklin, Bertie Vidgen, Amanpreet Singh, Douwe Kiela and Shikib Mehri	3303
<i>ThinkAnswer Loss: Balancing Semantic Similarity and Exact Matching for LLM Reasoning Enhancement</i>	
Shan Yang, Kun Wu, Zeju Li, Linlin Zhang, Xiangyu Pei, Leike An and Yu Liu	3325
<i>Detecting Stealthy Backdoor Samples based on Intra-class Distance for Large Language Models</i>	
Jinwen Chen, Hainan Zhang, Fei Sun, Qinnan Zhang, Sijia Wen, Ziwei Wang and Zhiming Zheng	3348
<i>Rust-doctor: Enhanced Feature for Rust Ownership and Lifetime Repair with Balanced Training Data Generation</i>	
Wenzhang Yang, Xiaoning Ren, Cuifeng Gao and Yinxing Xue	3366
<i>SLIM: Subtrajectory-Level Elimination for More Effective Reasoning</i>	
Xifeng Yao, Chengyuan Ma, Dongyu Lang, Yinhao Ni, Zhiwei Xu, Huarui Xie, Zihao Chen, Guang Shen, Dandan Tu, Yi Bai and Changzheng Zhang	3377
<i>From Cross-Task Examples to In-Task Prompts: A Graph-Based Pseudo-Labeling Framework for In-context Learning</i>	
Zihan Chen, Song Wang, Xingbo Fu, Chengshuai Shi, Zhenyu Lei, Cong Shen and Jundong Li	3395
<i>Instance-level Randomization: Toward More Stable LLM Evaluations</i>	
Yiyang Li, Yonghuang Wu, Ying Luo, Liangtai Sun, Zishu Qin, Lin Qiu, Xuezhi Cao and Xunliang Cai	3411
<i>Not All Voices Are Rewarded Equally: Probing and Repairing Reward Models across Human Diversity</i>	
Zihao Li, Feihao Fang, Xitong Zhang, Jiaru Zou, Zhining Liu, Wei Xiong, Ziwei Wu, Baoyu Jing and Jingrui He	3426
<i>PAMN: Multi-phase Correlation Modeling for Contrast-Enhanced 3D Medical Image Retrieval</i>	
Haonan Tong, Ke Liu, Chuang Zhang, Xinglin Zhang, Tao Chen, Jenq-Neng Hwang and Lei Li	3456
<i>Safety in Large Reasoning Models: A Survey</i>	
Cheng Wang, Yue Liu, Baolong Bi, Duzhen Zhang, Zhong-Zhi Li, Yingwei MA, Yufei He, Shengju Yu, Xinfeng Li, Junfeng Fang, Jiaheng Zhang and Bryan Hooi	3468
<i>SafeConf: A Confidence-Calibrated Safety Self-Evaluation Method for Large Language Models</i>	
Bo Zhang, Cong Gao, Linkang Yang, Bingxu Han, Minghao Hu, Zhunchen Luo, Guotong Geng, Xiaoying Bai, Jun Zhang, Wen Yao and Zhong Wang	3483

<i>DocAssistant: Integrating Key-region Reading and Step-wise Reasoning for Robust Document Visual Question Answering</i>	
Jinxu Zhang, Qiyuan Fan and Yu Zhang	3496
<i>LNE-Blocking: An Efficient Framework for Contamination Mitigation Evaluation on Large Language Models</i>	
Ruijie Hou, Yueyang Jiao, Hanxu Hu, Yingming Li, Wai Lam, Huajian Zhang and Hongyuan Lu	3512
<i>Enhancing Hate Speech Classifiers through a Gradient-assisted Counterfactual Text Generation Strategy</i>	
Michael Van Supranes, Shaowen Peng, Shoko Wakamiya and Eiji Aramaki	3529
<i>Learning SQL Like a Human: Structure-Aware Curriculum Learning for Text-to-SQL Generation</i>	
Xiaohu Zhu, Qian Li, Lizhen Cui and Yuntao Du	3545
<i>Chain-of-Interactions: Multi-step Iterative ICL Framework for Abstractive Task-Oriented Dialogue Summarization of Conversational AI Interactions</i>	
Jason S Lucas, Ali Al Lawati, Mahjabin Nahar, John Chen and Mahnoosh Mehrabani	3560
<i>Your Semantic-Independent Watermark is Fragile: A Semantic Perturbation Attack against EaaS Watermark</i>	
Zekun Fei, Biao Yi, Jianing Geng, He Ruiqi, Lihai Nie and Zheli Liu	3600
<i>Query Optimization for Parametric Knowledge Refinement in Retrieval-Augmented Large Language Models</i>	
Youan Cong, Pritom Saha Akash, Cheng Wang and Kevin Chen-Chuan Chang	3615
<i>SKA-Bench: A Fine-Grained Benchmark for Evaluating Structured Knowledge Understanding of LLMs</i>	
Zhiqiang Liu, Enpei Niu, Yin Hua, Mengshu Sun, Lei Liang, Huajun Chen and Wen Zhang	3626
<i>PD³F: A Pluggable and Dynamic DoS-Defense Framework against resource consumption attacks targeting Large Language Models</i>	
Yuanhe Zhang, Xinyue Wang, Haoran Gao, Zhenhong Zhou, Fanyu Meng, Yuyao Zhang and Sen Su	3641
<i>From Implicit Exploration to Structured Reasoning: Guideline and Refinement for LLMs</i>	
Jiaxiang Chen, Zhuo Wang, Mingxi Zou, Zhucong Li, Zhijian Zhou, Song Wang and Zenglin Xu	3672
<i>PIP: Perturbation-based Iterative Pruning for Large Language Models</i>	
Yi Cao, Wei-Jie Xu, Yucheng Shen, Weijie Shi, Chi-Min Chan, Jianfeng Qu and Jiajie Xu . .	3685
<i>Convolutional LoRA Aggregation for Unseen Tasks Adaptation</i>	
Xinhao Wu, Jialin Liu, Yutai Duan and Jie Liu	3702
<i>CDT: A Comprehensive Capability Framework for Large Language Models Across Cognition, Domain, and Task</i>	
Haosi Mo, Xinyu Ma, Xuebo Liu, Derek F. Wong, YU LI, Jie Liu and Min Zhang	3715
<i>Multilingual Collaborative Defense for Large Language Models</i>	
Hongliang Li, Jinan Xu, Gengping Cui, Changhao Guan, Fengran Mo and Kaiyu Huang . . .	3735
<i>Role-Guided Annotation and Prototype-Aligned Representation Learning for Historical Literature Sentiment Classification</i>	
Hongfei Du, Jiacheng Shi, Jacobo Myerston, Sidi Lu, Gang Zhou and Ashley Gao	3756

<i>MetaMixSpeech: Meta Task Augmentation for Low-Resource Speech Recognition</i>	
Yaqi Chen, Hao Zhang, Wenlin Zhang, XuKui Yang, Dan Qu and Yunpeng Liu	3769
<i>RECAST: Retrieval-Augmented Contextual ASR via Decoder-State Keyword Spotting</i>	
Ashish Mittal, Sunita Sarawagi and Preethi Jyothi	3780
<i>PREE: Towards Harmless and Adaptive Fingerprint Editing in Large Language Models via Knowledge Prefix Enhancement</i>	
Xubin Yue, Zhenhua Xu, Wenpeng Xing, Jiahui Yu, Mohan Li and Meng Han	3794
<i>Beyond Spurious Signals: Debiasing Multimodal Large Language Models via Counterfactual Inference and Adaptive Expert Routing</i>	
Zichen Wu, Hsiu-Yuan Huang and Yunfang Wu	3805
<i>Text-centric Alignment for Bridging Test-time Unseen Modality</i>	
Yun-Da Tsai, Ting-Yu Yen, Pei-Fu Guo, Zhe-Yan Li and Shou-De Lin	3826
<i>HierPrompt: Zero-Shot Hierarchical Text Classification with LLM-Enhanced Prototypes</i>	
Qian Zhang, Qinliang Su, Wei Zhu and Pang Yachun	3846
<i>RouterEval: A Comprehensive Benchmark for Routing LLMs to Explore Model-level Scaling Up in LLMs</i>	
Zhongzhan Huang, Guoming Ling, Yupei Lin, Yandong Chen, Shanshan Zhong, Hefeng Wu and Liang Lin	3860
<i>Can We Steer Reasoning Direction by Thinking Intervention?</i>	
Kingsheng Zhang, Luxi Xing, Chen Zhang, Yanbing Liu, Yifan Deng, Yunpeng Li, Yue Hu and Chenxu Niu	3888
<i>MPO: Boosting LLM Agents with Meta Plan Optimization</i>	
Weimin Xiong, Yifan Song, Qingxiu Dong, Bingchan Zhao, Feifan Song, XWang and Sujian Li	3914
<i>Exploring the Generalizability of Factual Hallucination Mitigation via Enhancing Precise Knowledge Utilization</i>	
Si yuan Zhang, Yichi Zhang, Yinpeng Dong and Hang Su	3936
<i>Learning What to Remember: Adaptive Probabilistic Memory Retention for Memory-Efficient Language Models</i>	
S M Rafiuddin and Muntaha Nujat Khan	3969
<i>Unlocking Smarter Device Control: Foresighted Planning with a World Model-Driven Code Execution Approach</i>	
Xiaoran Yin, Xu Luo, Hao Wu, Lianli Gao and Jingkuan Song	3982
<i>RGAR: Recurrence Generation-augmented Retrieval for Factual-aware Medical Question Answering</i>	
Sichu Liang, Linhai Zhang, Hongyu Zhu, Wenwen Wang, Yulan He and Deyu Zhou	4006
<i>EcoSafeRAG: Efficient Security through Context Analysis in Retrieval-Augmented Generation</i>	
Ruobing Yao, Yifei Zhang, Shuang Song, Neng Gao and Chenyang Tu	4034
<i>StereoDetect: Detecting Stereotypes and Anti-stereotypes the Correct Way Using Social Psychological Underpinnings</i>	
Kaustubh Shivshankar Shejole and Pushpak Bhattacharyya	4051

<i>Sparkle: Mastering Basic Spatial Capabilities in Vision Language Models Elicits Generalization to Spatial Reasoning</i>	
Yihong Tang, Ao Qu, Zhaokai Wang, Dingyi Zhuang, Zhaofeng Wu, Wei Ma, Shenhao Wang, Yunhan Zheng, Zhan Zhao and Jinhua Zhao	4083
<i>How Does Knowledge Selection Help Retrieval Augmented Generation?</i>	
Xiangci Li and Jessica Ouyang	4104
<i>UPLex: Fine-Grained Personality Control in Large Language Models via Unsupervised Lexical Modulation</i>	
Tianlong Li, Wenhao Liu, Muling Wu, Shihan Dou, Zhenghua Wang, Changze Lv, Xiaohua Wang, Xiaoqing Zheng and Xuanjing Huang	4122
<i>ParetoRAG: Leveraging Sentence-Context Attention for Robust and Efficient Retrieval-Augmented Generation</i>	
Ruobing Yao, Yifei Zhang, Shuang Song, Yuhan Liu, Neng Gao and Chenyang Tu	4137
<i>FlexQuant: A Flexible and Efficient Dynamic Precision Switching Framework for LLM Quantization</i>	
Fangxin Liu, Zongwu Wang, Jinhong Xia, Junping Zhao, Shouren Zhao, Jinjin Li, Jian Liu, Li Jiang and Haibing Guan	4152
<i>ReLoop: Seeing Twice and Thinking Backwards via Closed-loop Training to Mitigate Hallucinations in Multimodal understanding</i>	
Jianjiang Yang, Yanshu Li and Ziyang Huang	4162
<i>Sequence Structure Aware Retriever for Procedural Document Retrieval: A New Dataset and Baseline</i>	
Zhenqi Ye, HaoPeng Ren, Yi Cai, Qingbao Huang, Jing Qin, Pinli Zhu and Songwen Gong	4180
<i>The Effect of Language Diversity When Fine-Tuning Large Language Models for Translation</i>	
David Stap and Christof Monz	4199
<i>David vs. Goliath: Cost-Efficient Financial QA via Cascaded Multi-Agent Reasoning</i>	
Chenghao Liu, Qian Liu, Ziqin Zhu, Hao Fei and Aniket Mahanti	4212
<i>Benchmarking Uncertainty Metrics for LLM Target-Aware Search</i>	
Pei-Fu Guo, Yun-Da Tsai and Shou-De Lin	4230
<i>ZOGRASCOPE: A New Benchmark for Semantic Parsing over Property Graphs</i>	
Francesco Cazzaro, Justin Kleindienst, Sofia Márquez Gomez and Ariadna Quattoni	4239
<i>FG-PRM: Fine-grained Hallucination Detection and Mitigation in Language Model Mathematical Reasoning</i>	
Ruosan Li, Ziming Luo and Xinya Du	4247
<i>Recipe2Plan: Evaluating Planning Abilities of LLMs for Efficient and Feasible Multitasking with Time Constraints Between Actions</i>	
Zirui Wu, Xiao Liu, Jiayi Li, Lingpeng Kong and Yansong Feng	4279
<i>Unlocking the Effectiveness of LoRA-FP for Seamless Transfer Implantation of Fingerprints in Downstream Models</i>	
Zhenhua Xu, Zhaokun Yan, Binhan Xu, Xin Tong, Haitao Xu, Yourong Chen and Meng Han	4302
<i>AELC: Adaptive Entity Linking with LLM-Driven Contextualization</i>	
Fang Wang, Zhengwei Tao, Ming Wang, Minghao Hu and Xiaoying Bai	4313
<i>MetaLadder: Ascending Mathematical Solution Quality via Analogical-Problem Reasoning Transfer</i>	
Honglin Lin, Zhuoshi Pan, Qizhi Pei, Xin Gao, Yu Li, Mengzhang Cai, Conghui He and Lijun Wu	4328

<i>GLProtein: Global-and-Local Structure Aware Protein Representation Learning</i> Yunqing Liu, Wenqi Fan, Xiaoyong Wei and Li Qing	4355
<i>Reward Mixology: Crafting Hybrid Signals for Reinforcement Learning Driven In-Context Learning</i> Changshuo Zhang, Ang Gao, Xiao Zhang, Yong Liu, Deyang Li, Fangchao Liu and Xinyu Zhang 4373	
<i>Can Multimodal LLMs See Materials Clearly? A Multimodal Benchmark on Materials Characterization</i> Zhengzhao Lai, Youbin Zheng, Zhenyang Cai, Haonan Lyu, Jingpu Yang, Hong-Qing Liang, Yan Hu and Benyou Wang	4384
<i>GRADE: Generating multi-hop QA and fine-grained Difficulty matrix for RAG Evaluation</i> Jeongsoo Lee, Daeyong Kwon and Kyohoon Jin	4405
<i>FusionDTI: Fine-grained Binding Discovery with Token-level Fusion for Drug-Target Interaction</i> Zhaohan Meng, Zaiqiao Meng, Ke Yuan and Iadh Ounis	4425
<i>A Survey on Training-free Alignment of Large Language Models</i> Birong Pan, Yongqi Li, Weiyu Zhang, Wenpeng Lu, Mayi Xu, Shen Zhou, Yuanyuan Zhu, Ming Zhong and Tiejun Qian	4445
<i>CIVET: Systematic Evaluation of Understanding in VLMs</i> Massimo Rizzoli, Simone Alghisi, Olha Khomyn, Gabriel Roccabruna, Seyed Mahed Mousavi and Giuseppe Riccardi	4462
<i>How Does Cognitive Bias Affect Large Language Models? A Case Study on the Anchoring Effect in Price Negotiation Simulations</i> Yoshiki Takenami, Yin Jou Huang, Yugo Murawaki and Chenhui Chu	4481
<i>Enhancing Speech-to-Speech Dialogue Modeling with End-to-End Retrieval-Augmented Generation</i> Pengchao Feng, Ziyang Ma, Wenxi Chen, Yao Li, Sheng Wang, Kai Yu and Xie Chen	4499
<i>Backdoor-Powered Prompt Injection Attacks Nullify Defense Methods</i> Yulin Chen, Haoran Li, Yuan Sui, Yangqiu Song and Bryan Hooi	4508
<i>Path-enhanced Pre-trained Language Model for Knowledge Graph Completion</i> Hao Wang, Dandan Song, Zhijing Wu, Yuhang Tian and Pan Yang	4528
<i>Zero-shot Cross-lingual NER via Mitigating Language Difference: An Entity-aligned Translation Perspective</i> Zhihao Zhang, Sophia Yat Mei Lee, Dong Zhang, Shoushan Li and Guodong Zhou	4541
<i>Zero-Shot Cross-Domain Aspect-Based Sentiment Analysis via Domain-Contextualized Chain-of-Thought Reasoning</i> Chuming Shen, Wei Wei, Dong Wang and Zhong-Hao Wang	4558
<i>Tree of Agents: Improving Long-Context Capabilities of Large Language Models through Multi-Perspective Reasoning</i> Song Yu, Xiaofei Xu, KE Deng, Li Li and Lin Tian	4574
<i>Cross-Cultural Transfer of Commonsense Reasoning in LLMs: Evidence from the Arab World</i> Saeed Almheiri, Rania Elbadry, Mena Attia, Chenxi Wang, Preslav Nakov, Timothy Baldwin and Fajri Koto	4593
<i>Enhancing Partially Relevant Video Retrieval with Robust Alignment Learning</i> Long Zhang, Peipei Song, Jianfeng Dong, Kun Li and Xun Yang	4615

<i>Multi-level Diagnosis and Evaluation for Robust Tabular Feature Engineering with Large Language Models</i>	
Yebin Lim and Susik Yoon	4630
<i>Prejudge-Before-Think: Enhancing Large Language Models at Test-Time by Process Prejudge Reasoning</i>	
Jianing Wang, Jin Jiang, Yang Liu, Mengdi Zhang and Xunliang Cai	4656
<i>From: Frobenius Norm-Based Data-Free Adaptive Model Merging</i>	
Zijian Li, Xiaocheng Feng, Huixin Liu, Yichong Huang, Ting Liu and Bing Qin	4674
<i>Dynamic Simulation Framework for Disinformation Dissemination and Correction With Social Bots</i>	
Boyu Qiao, Kun Li, Wei Zhou and Songlin Hu	4688
<i>Beyond the First Error: Process Reward Models for Reflective Mathematical Reasoning</i>	
Zhaohui Yang, Chenghua He, Xiaowen Shi, Shihong Deng, Linjing Li, Qiyue Yin and Daxin Jiang	4711
<i>PrAd: Prompt Adaptive Tuning for Decoder-only Language Models</i>	
Youneng Ma, Junyi He and Haojun Fei	4729
<i>Personalized Question Answering with User Profile Generation and Compression</i>	
Hang Su, Yun Yang, Tianyang Liu, Xin Liu, Peng Pu and Xuesong Lu	4744
<i>Dream to Chat: Model-based Reinforcement Learning on Dialogues with User Belief Modeling</i>	
Yue Zhao, Xiaoyu Wang, Dan Wang, Zhonglin Jiang, Qingqing Gu, Teng Chen, Ningyuan Xi, Jinxian Qu, Yong Chen and Luo Ji	4764
<i>FakeSV-VLM: Taming VLM for Detecting Fake Short-Video News via Progressive Mixture-Of-Experts Adapter</i>	
JunXi Wang, Yaxiong Wang, Lechao Cheng and Zhun Zhong	4782
<i>Beyond Inherent Cognition Biases in LLM-Based Event Forecasting: A Multi-Cognition Agentic Framework</i>	
Zhen Wang, Xi Zhou, Yating Yang, Bo Ma, Lei Wang, Rui Dong and Azmat Anwar	4799
<i>Breaking the Reviewer: Assessing the Vulnerability of Large Language Models in Automated Peer Review Under Textual Adversarial Attacks</i>	
Tzu-Ling Lin, Wei-Chih Chen, Teng-Fang Hsiao, Hou-I Liu, Ya-Hsin Yeh, Yu-Kai Chan, Wen-Sheng Lien, Po-Yen Kuo, Philip S. Yu and Hong-Han Shuai	4819
<i>Watermarking with Low-Entropy POS-Guided Token Partitioning and Z-Score-Driven Dynamic Bias for Large Language Models</i>	
He Li, Xiaojun Chen, Zhendong Zhao, Yunfei Yang, Xin Zhao and Jingcheng He	4840
<i>Knowledge Graph-Driven Memory Editing with Directional Interventions</i>	
Jinhu Fu, Kun Wang, Chongye Guo, Junfeng Fang, Wentao Zhang and Sen Su	4860
<i>DTDES-KGE: Dual-Teacher Knowledge Distillation with Distinct Embedding Spaces for Knowledge Graph Embeddings</i>	
Bofan Wei, Hongyuan Xu, Yuhang Niu, Jiarui Ren, Yanlong Wen and Xiaojie Yuan	4875
<i>LLMEval-Med: A Real-world Clinical Benchmark for Medical LLMs with Physician Validation</i>	
Ming Zhang, Yujiong Shen, Zelin Li, Huayu Sha, Binze Hu, Yuhui Wang, Chenhao Huang, Shichun Liu, Jingqi Tong, Changhao Jiang, Mingxu Chai, Zhiheng Xi, Shihan Dou, Tao Gui, Qi Zhang and Xuanjing Huang	4888

<i>Watermark Smoothing Attacks against Language Models</i>	
Hongyan Chang, Hamed Hassani and Reza Shokri	4915
<i>PICD-Instruct: A Generative Instruction Learning Framework for Few-Shot Multi-Intent Spoken Language Understanding</i>	
Wenbin Hua, Rui Fan, Tingting He and Ming Dong	4942
<i>Forewarned is Forearmed: Pre-Synthesizing Jailbreak-like Instructions to Enhance LLM Safety Guardrail to Potential Attacks</i>	
Sheng Liu, Qiang Sheng, Danding Wang, Yang Li, Guang Yang and Juan Cao	4957
<i>Are Knowledge and Reference in Multilingual Language Models Cross-Lingually Consistent?</i>	
Xi Ai, Mahardika Krisna Ihsani and Min-Yen Kan	4975
<i>Krikri: Advancing Open Large Language Models for Greek</i>	
Dimitris Roussis, Leon Voukoutis, Georgios Paraskevopoulos, Sokratis Sofianopoulos, Prokopis Prokopidis, Vassilis Papavassileiou, Athanasios Katsamanis, Stelios Piperidis and Vassilis Katsourou	5012
<i>Beyond the Scientific Document: A Citation-Aware Multi-Granular Summarization Approach with Heterogeneous Graphs</i>	
Quoc-An Nguyen, Xuan-Hung Le, Thi-Minh-Thu Vu and Hoang-Quynh Le	5034
<i>Detecting Continuously Evolving Scam Calls under Limited Annotation: A LLM-Augmented Expert Rule Framework</i>	
Haoyu Ma, Qinliang Su, Minhua Huang and Wu Kai	5047
<i>An Empirical Study of Position Bias in Modern Information Retrieval</i>	
Ziyang Zeng, Dun Zhang, Jiacheng Li, Zoupanxiang, Yudong Zhou and Yuqing Yang	5069
<i>GenPoE: Generative Passage-level Mixture of Experts for Knowledge Enhancement of LLMs</i>	
Xuebing Liu, Shanbao Qiao and Seung-Hoon Na	5082
<i>CoRanking: Collaborative Ranking with Small and Large Ranking Agents</i>	
Wenhan Liu, Xinyu Ma, Yutao Zhu, Lixin Su, Shuaiqiang Wang, Dawei Yin and Zhicheng Dou	5098
<i>HIRAG: Hierarchical-Thought Instruction-Tuning Retrieval-Augmented Generation</i>	
Yihan Jiao, Zhehao Tan, Dan Yang, Duolin Sun, Jie Feng, Yue Shen, Jian Wang and Peng Wei	5111
<i>Towards Personalized Conversational Sales Agents: Contextual User Profiling for Strategic Action</i>	
Tongyoung Kim, Jeongeun Lee, SooJin Yoon, SungHwan Kim and Dongha Lee	5131
<i>WebCoT: Enhancing Web Agent Reasoning by Reconstructing Chain-of-Thought in Reflection, Branching, and Rollback</i>	
Minda Hu, Tianqing Fang, Jianshu Zhang, Jun-Yu Ma, Zhisong Zhang, Jingyan Zhou, Hongming Zhang, Haitao Mi, Dong Yu and Irwin King	5155
<i>Interesting Culture: Social Relation Recognition from Videos via Culture De-confounding</i>	
Yuxuan Zhang, Yangfu Zhu, Haorui Wang and Bin Wu	5174
<i>ThinkSwitcher: When to Think Hard, When to Think Fast</i>	
Guosheng Liang, Longguang Zhong, Ziyi Yang and Xiaojun Quan	5185
<i>MaGiX: A Multi-Granular Adaptive Graph Intelligence Framework for Enhancing Cross-Lingual RAG</i>	
Nguyen Manh Hieu, Vu Lam Anh, Hung Pham Van, Nam Le Hai, Linh Ngo Van, Nguyen Thi Ngoc Diep and Thien Huu Nguyen	5202

<i>LexTime: A Benchmark for Temporal Ordering of Legal Events</i>	
Claire Barale, Leslie Barrett, Vikram Sunil Bajaj and Michael Rovatsos	5220
<i>Beyond the Surface: A Solution-Aware Retrieval Model for Competition-level Code Generation</i>	
Shiwen Zhang, Lingxiang Wang, Hainan Zhang, Ziwei Wang, Sijia Wen and Zhiming Zheng	5237
<i>X-Boundary: Establishing Exact Safety Boundary to Shield LLMs from Jailbreak Attacks without Compromising Usability</i>	
Xiaoya Lu, Dongrui Liu, Yi Yu, Luxin Xu and Jing Shao	5247
<i>Tag&Tab: Pretraining Data Detection in Large Language Models Using Keyword-Based Membership Inference Attack</i>	
Sagiv Antebi, Edan Habler, Asaf Shabtai and Yuval Elovici	5273
<i>EcoLANG: Efficient and Effective Agent Communication Language Induction for Social Simulation</i>	
Xinyi Mou, Chen Qian, Wei Liu, Ling Yan, Yao Hu, Xuanjing Huang and Zhongyu Wei ...	5287
<i>Revealing the Inherent Instructability of Pre-Trained Language Models</i>	
Seokhyun An, Minji Kim and Hyoungun Kim	5305
<i>What Media Frames Reveal About Stance: A Dataset and Study about Memes in Climate Change Discourse</i>	
Shijia Zhou, Siyao Peng, Simon M. Luebke, Jörg Haßler, Mario Haim, Saif M. Mohammad and Barbara Plank	5337
<i>Rethinking Personality Assessment from Human-Agent Dialogues: Fewer Rounds May Be Better Than More</i>	
Baiqiao Zhang, Zhifeng Liao, Xiangxian Li, Chao Zhou, Juan Liu, Xiaojuan Ma and Yulong Bian	5357
<i>TailorRPA: A Retrieval-Based Framework for Eliciting Personalized and Coherent Role-Playing Agents in General Domain</i>	
Zhenpeng Gao, Xiaofen Xing and Xiangmin Xu	5381
<i>SCE: Semantic Consistency Enhanced Reinforcement Learning for Multi-Hop Knowledge Graph Reasoning</i>	
Yanwen Huang, Yao Liu, Qiao Liu, Rui Hou and Tingting Dai	5413
<i>ReGraphRAG: Reorganizing Fragmented Knowledge Graphs for Multi-Perspective Retrieval-Augmented Generation</i>	
Soohyeong Kim, Seok Jun Hwang, JungHyouon Kim, Jeonghyeon Park and Yong Suk Choi .	5426
<i>GASE: Generatively Augmented Sentence Encoding</i>	
Manuel Frank and Haithem Afli	5444
<i>The rin womanstands for rights. Auditing LLMs in Uncovering Social Dynamics in Implicit Misogyny</i>	
Arianna Muti, Chris Emmery, Debora Nozza, Alberto Barrón-Cedeño and Tommaso Caselli	5462
<i>Fact Verification on Knowledge Graph via Programmatic Graph Reasoning</i>	
Yuanzhen Hao and Desheng Wu	5480
<i>Agent Trading Arena: A Study on Numerical Understanding in LLM-Based Agents</i>	
Tianmi Ma, Jiawei Du, Wenxin Huang, Wenjie Wang, Liang Xie, Xian Zhong and Joey Tianyi Zhou	5496
<i>Why We Feel What We Feel: Joint Detection of Emotions and Their Opinion Triggers in E-commerce</i>	
Arnav Attri, Anuj Attri, Suman Banerjee, Amey Patil, Muthusamy Chelliah, Nikesh Garera and Pushpak Bhattacharyya	5515

<i>Use Random Selection for Now: Investigation of Few-Shot Selection Strategies in LLM-based Text Augmentation</i>	
Jan Cegin, Branislav Pecher, Jakub Simko, Ivan Srba, Maria Bielikova and Peter Brusilovsky	5533
<i>BanglaByT5: Byte-Level Modelling for Bangla</i>	
Pramit Bhattacharyya and Arnab Bhattacharya	5551
<i>XTRA: Cross-Lingual Topic Modeling with Topic and Representation Alignments</i>	
Tien Phat Nguyen, Ngo Vu Minh, Tung Nguyen, Linh Ngo Van, Duc Anh Nguyen, Dinh Viet Sang and Trung Le	5561
<i>CodeContests+: High-Quality Test Case Generation for Competitive Programming</i>	
Zihan Wang, Siyao Liu, Yang Sun, Ming Ding and Hongyan Li	5576
<i>SPO: Self Preference Optimization with Self Regularization</i>	
Yuhao Sun, Yifan Zhang, Quandong Wang, Qinzhuo Wu, Wei Liu and Jian Luan	5601
<i>Long-context Language Models Fail in Basic Retrieval Tasks Without Sufficient Reasoning Steps</i>	
Yijiong Yu, Zhixiao Qi, Yongfeng Huang, Wei Wang, Weifeng.liu, Ran Chen and Ji Pei	5615
<i>Benchmarking Critical Questions Generation: A Challenging Reasoning Task for Large Language Models</i>	
Blanca Calvo Figueras and Rodrigo Agerri	5635
<i>ResearchArena: Benchmarking Large Language Models' Ability to Collect and Organize Information as Research Agents</i>	
Hao Kang and Chenyan Xiong	5653
<i>LLMs are Privacy Erasable</i>	
Zipeng Ye and Wenjian Luo	5672
<i>How Good are LLM-based Rerankers? An Empirical Analysis of State-of-the-Art Reranking Models</i>	
Abdelrahman Abdallah, Bhawna Piryani, Jamshid Mozafari, Mohammed Ali and Adam Jatowt	5693
<i>DeAR: Dual-Stage Document Reranking with Reasoning Agents via LLM Distillation</i>	
Abdelrahman Abdallah, Jamshid Mozafari, Bhawna Piryani and Adam Jatowt	5710
<i>CANDY: Benchmarking LLMs' Limitations and Assistive Potential in Chinese Misinformation Fact-Checking</i>	
Ruiling Guo, Xinwei Yang, Chen Huang, Tong Zhang and Yong Hu	5724
<i>E-Verify: A Paradigm Shift to Scalable Embedding-based Factuality Verification</i>	
Zeyang Liu, Jingfeng Xue, Xiuqi Yang, Wenbiao Du, Jiarun Fu, Junbao Chen, Wenjie Guo and Yong Wang	5759
<i>LLM Jailbreak Detection for (Almost) Free!</i>	
Guorui Chen, Yifan Xia, Xiaojun Jia, Zhijiang Li, Philip Torr and Jindong Gu	5777
<i>When to Continue Thinking: Adaptive Thinking Mode Switching for Efficient Reasoning</i>	
Xiaoyun Zhang, Jingqing Ruan, Xing Ma, Yawen Zhu, Haodong Zhao, Hao Li, Jiansong Chen, Ke Zeng and Xunliang Cai	5808
<i>Plugging Schema Graph into Multi-Table QA: A Human-Guided Framework for Reducing LLM Reliance</i>	
Xixi Wang, Miguel Costa, Jordanka Kovaceva, Shuai Wang and Francisco C. Pereira	5829

<i>Evolution in Simulation: AI-Agent School with Dual Memory for High-Fidelity Educational Dynamics</i> Sheng Jin, Haoming Wang, Zhiqi Gao, Yongbo Yang, Bao Chunjia and Chengliang Wang . .	5843
<i>Retrieval-Augmented Machine Translation with Unstructured Knowledge</i> Jiaan Wang, Fandong Meng, Yingxue Zhang and Jie Zhou	5858
<i>MARS-Bench: A Multi-turn Athletic Real-world Scenario Benchmark for Dialogue Evaluation</i> Chenghao Yang, Yinbo Luo, Zhoufutu Wen, Qi Chu, Tao Gong, Longxiang Liu, Kaiyuan Zhang, Jianpeng Jiao, Ge Zhang, Wenhao Huang and Nenghai Yu	5872
<i>UTMath: A Benchmark for Math Evaluation with Unit Test</i> Bo Yang, Qingping Yang, Yingwei MA and Runtao Liu	5899
<i>The Green KNIGHT: Green Machine Translation with Knowledge-Distilled, Narrow, Inexpensive, Greedy, Hybrid Transformers</i> Andreas Guta, Frithjof Petrick and Peter Polák	5916
<i>Constructing Your Model's Value Distinction: Towards LLM Alignment with Anchor Words Tuning</i> Zhen Yang, Ping Jian, Chengzhi Li, Chenxu Wang, Xinyue Zhang and Wenpeng Lu	5932
<i>MCiteBench: A Multimodal Benchmark for Generating Text with Citations</i> Caiyu Hu, Yikai Zhang, Tinghui Zhu, Yiwei Ye and Yanghua Xiao	5949
<i>Do LLMs Know and Understand Domain Conceptual Knowledge?</i> Sijia Shen, Feiyan Jiang, Peiyan Wang, Yubo Feng, Yuchen Jiang and Chang Liu	5967
<i>Agent Laboratory: Using LLM Agents as Research Assistants</i> Samuel Schmidgall, Yusheng Su, Ze Wang, Ximeng Sun, Jialian Wu, Xiaodong Yu, Jiang Liu, Michael Moor, Zicheng Liu and Emad Barsoum	5977
<i>Retrieval-Augmented Generation with Hierarchical Knowledge</i> Haoyu Huang, Yongfeng Huang, Yang Junjie, Zhenyu Pan, Yongqiang Chen, Kaili Ma, Hongzhi Chen and James Cheng	6044
<i>Regularized Contrastive Decoding with Hard Negative Samples for LLM Hallucination Mitigation</i> Haonan Sheng, Dou Hu, Lingwei Wei, Wei Zhou and Songlin Hu	6061
<i>CharacterCraft: Bridging the Literature-Reality Dialogue Gap for Practical Role-Playing Agents</i> Xuyan Yin, Xinran Yang, Zihao Li, Lixin Zou and Chenliang Li	6074
<i>Drift: Decoding-time Personalized Alignments with Implicit User Preferences</i> Minbeom Kim, Kang-il Lee, Seongho Joo, Hwaran Lee, Thibaut Thonet and Kyomin Jung .	6107
<i>Discovering Semantic Subdimensions through Disentangled Conceptual Representations</i> Yunhao Zhang, Shaonan Wang, Nan Lin, Xinyi Dong, Chong Li and Chengqing Zong	6127
<i>Identifying Aspects in Peer Reviews</i> Sheng Lu, Iliia Kuznetsov and Iryna Gurevych	6145
<i>Tree-Structured Non-Autoregressive Decoding for Sequence-to-Sequence Text Generation</i> Pengyu Ji, Yufei Liu, Xiang Hu and Kewei Tu	6168
<i>Towards More Efficient Post-training via Fourier Domain Adapter Framework</i> Yijia Fan, Jusheng Zhang and Keze Wang	6175
<i>KERAG: Knowledge-Enhanced Retrieval-Augmented Generation for Advanced Question Answering</i> Yushi Sun, Kai Sun, Yifan Ethan Xu, Xiao Yang, Xin Luna Dong, Nan Tang and Lei Chen .	6194

<i>Not All Features Deserve Attention: Graph-Guided Dependency Learning for Tabular Data Generation with Language Models</i>	
Zheyu Zhang, Shuo Yang, Bardh Prenkaj and Gjergji Kasneci	6217
<i>CCG: Rare-Label Prediction via Neural SEM-Driven Causal Game</i>	
Yijia Fan, Jusheng Zhang, Kaitong Cai, Jing Yang and Keze Wang	6243
<i>Multimodal Emotion Recognition in Conversations: A Survey of Methods, Trends, Challenges and Prospects</i>	
ChengYan Wu, Yiqiang Cai, Yang Liu, Pengxu Zhu, Yun Xue, Ziwei Gong, Julia Hirschberg and Bolei Ma	6257
<i>When Allies Turn Foes: Exploring Group Characteristics of LLM-Based Multi-Agent Collaborative Systems Under Adversarial Attacks</i>	
Jiahao Zhang, Baoshuo Kan, Tao Gong, Fu Lee Wang and Tianyong Hao	6275
<i>EditID: Training-Free Editable ID Customization for Text-to-Image Generation</i>	
Guandong Li and Zhaobin Chu	6301
<i>OSC: Cognitive Orchestration through Dynamic Knowledge Alignment in Multi-Agent LLM Collaboration</i>	
Jusheng Zhang, Yijia Fan, Kaitong Cai, Xiaofei Sun and Keze Wang	6320
<i>VideoLLM Knows When to Speak: Enhancing Time-Sensitive Video Comprehension with Video-Text Duet Interaction Format</i>	
Yueqian Wang, Xiaojun Meng, Yuxuan Wang, Jianxin Liang, Jiansheng Wei, Huishuai Zhang and Dongyan Zhao	6338
<i>To Answer or Not to Answer (TAONA): A Robust Textual Graph Understanding and Question Answering Approach</i>	
Yuchen Yan, Aakash Kolekar, Sahika Genc, Wenju Xu, Edward W Huang, Anirudh Srinivasan, Mukesh Jain, Qi He and Hanghang Tong	6360
<i>Understanding Refusal in Language Models with Sparse Autoencoders</i>	
Wei Jie Yeo, Nirmalendu Prakash, Clement Neo, Ranjan Satapathy, Roy Ka-Wei Lee and Erik Cambria	6377
<i>Where Did That Come From? Sentence-Level Error-Tolerant Attribution</i>	
Ori Ernst, Aviv Slobodkin, Meng Cao, Sihui Wei and Jackie CK Cheung	6400
<i>Alleviating Performance Degradation Caused by Out-of-Distribution Issues in Embedding-Based Retrieval</i>	
Haotong Bao, Jianjin Zhang, Qi Chen, Weihao Han, Zhengxin Zeng, Ruiheng Chang, Mingzheng Li, Hao Sun, Weiwei Deng, Feng Sun and Qi Zhang	6418
<i>Can LLMs Find a Needle in a Haystack? A Look at Anomaly Detection Language Modeling</i>	
Leslie Barrett, Vikram Sunil Bajaj and Robert John Kingan	6428
<i>Beyond Single Frames: Can LMMs Comprehend Implicit Narratives in Comic Strip?</i>	
Xiaochen Wang, Heming Xia, Jialin Song, Longyu Guan, Qingxiu Dong, Rui Li, Yixin Yang, Yifan Pu, Weiyao Luo, Yiru Wang, Xiangdi Meng, Wenjie Li and Zhifang Sui	6436
<i>Enhancing Multi-Agent Debate System Performance via Confidence Expression</i>	
Zijie Lin and Bryan Hooi	6453
<i>The Face of Persuasion: Analyzing Bias and Generating Culture-Aware Ads</i>	
Aysan Aghazadeh and Adriana Kovashka	6472

<i>SIFT: Grounding LLM Reasoning in Contexts via Stickers</i>	
Zihao Zeng, Xuyao Huang, Boxiu Li and Zhijie Deng	6501
<i>When Inverse Data Outperforms: Exploring the Pitfalls of Mixed Data in Multi-Stage Fine-Tuning</i>	
Mengyi Deng, Xin Li, Tingyu Zhu, Zhicheng Yang, Zhijiang Guo and Wei Wang	6514
<i>LUME: LLM Unlearning with Multitask Evaluations</i>	
Anil Ramakrishna, Yixin Wan, Xiaomeng Jin, Kai-Wei Chang, Zhiqi Bu, Bhanukiran Vinzamuri, Volkan Cevher, Mingyi Hong and Rahul Gupta	6524
<i>How do Language Models Generate Slang: A Systematic Comparison between Human and Machine-Generated Slang Usages</i>	
Siyang Wu and Zhewei Sun	6536
<i>Bridging the Dynamic Perception Gap: Training-Free Draft Chain-of-Thought for Dynamic Multimodal Spatial Reasoning</i>	
Siqu Ou, Hongcheng Liu, Pingjie Wang, Yusheng Liao, Chuan Xuan, Yanfeng Wang and Yu Wang	6560
<i>MedCOD: Enhancing English-to-Spanish Medical Translation of Large Language Models Using Enriched Chain-of-Dictionary Framework</i>	
Md Shahidul Salim, Lian Fu, Arav Adikesh Ramakrishnan, Zonghai Yao and Hong yu	6579
<i>Chatbot To Help Patients Understand Their Health</i>	
Won Seok Jang, Hieu Tran, Manav Shaileshkumar Mistry, Sai Kiran Gandluri, Yifan Zhang, Sharmin Sultana, Sunjae Kwon, Yuan Zhang, Zonghai Yao and Hong yu	6598
<i>A Knapsack by Any Other Name: Presentation impacts LLM performance on NP-hard problems</i>	
Alex Duchnowski, Ellie Pavlick and Alexander Koller	6628
<i>Is Safety Standard Same for Everyone? User-Specific Safety Evaluation of Large Language Models</i>	
Yeonjun In, Wonjoong Kim, Kanghoon Yoon, Sungchul Kim, Mehrab Tanjim, Sangwu Park, Kibum Kim and Chanyoung Park	6652
<i>Jailbreak Attack Initializations as Extractors of Compliance Directions</i>	
Amit LeVi, Rom Himmelstein, Yaniv Nemcovsky, Avi Mendelson and Chaim Baskin	6672
<i>Train Once for All: A Transitional Approach for Efficient Aspect Sentiment Triplet Extraction</i>	
Xinmeng Hou, Lingyue Fu, Chenhao Meng, Kounianhua Du and Hai Hu	6706
<i>A Comprehensive Survey on the Trustworthiness of Large Language Models in Healthcare</i>	
Manar Aljohani, Jun Hou, Sindhura Kommu and Xuan Wang	6720
<i>Self-Correction Makes LLMs Better Parsers</i>	
Ziyan Zhang, Yang Hou, Chen Gong and Zhenghua Li	6749
<i>Explaining Length Bias in LLM-Based Preference Evaluations</i>	
Zhengyu Hu, Linxin Song, Jieyu Zhang, Zheyuan Xiao, Tianfu Wang, Zhengyu Chen, Nicholas Jing Yuan, Jianxun Lian, Kaize Ding and Hui Xiong	6763
<i>Investigating Controversy Framing across Topics on Social Media</i>	
Maxwell Weinzierl and Sanda M. Harabagiu	6795
<i>HEAL: Hybrid Enhancement with LLM-based Agents for Text-attributed Hypergraph Self-supervised Representation Learning</i>	
Ruochang Li, Xiao Luo, Zhiping Xiao, Wei Ju and Ming Zhang	6815

<i>ReMamba: Equip Mamba with Effective Long-Sequence Modeling</i>	
Danlong Yuan, Jiahao Liu, Bei Li, Huishuai Zhang, Jingang Wang, Xunliang Cai and Dongyan Zhao	6830
<i>QUITO-X: A New Perspective on Context Compression from the Information Bottleneck Theory</i>	
Yihang Wang, Xu Huang, Bowen Tian, Yueyang Su, Lei Yu, Huaming Liao, Yixing Fan, Jiafeng Guo and Xueqi Cheng	6841
<i>Conv-Basis: A New Paradigm for Efficient Attention Inference and Gradient Computation in Transformers</i>	
Yingyu Liang, Heshan Liu, Zhenmei Shi, Zhao Song, Zhuoyan Xu, Jiale Zhao and Zhen Zhuang	6857
<i>Mitigating Gender Bias via Fostering Exploratory Thinking in LLMs</i>	
Kangda Wei, Hasnat Md Abdullah and Ruihong Huang	6895
<i>Beyond the Textual: Generating Coherent Visual Options for MCQs</i>	
Wanqiang Wang, Longzhu He and Wei Zheng	6918
<i>SafeSwitch: Steering Unsafe LLM Behavior via Internal Activation Signals</i>	
Peixuan Han, Cheng Qian, Xiushi Chen, Yuji Zhang, Heng Ji and Denghui Zhang	6936
<i>MADD: Multi-Agent Drug Discovery Orchestra</i>	
Gleb Vitalevich Solovev, Alina Borisovna Zhidkovskaya, Anastasia Orlova, Nina Gubina, Anastasia Vepreva, Rodion Golovinskii, Ilya Tonkii, Ivan Dubrovsky, Ivan Gurev, Dmitry Gilemkhanov, Denis Chistiakov, Timur A. Aliev, Ivan Poddiakov, Galina Zubkova, Ekaterina V. Skorb, Vladimir Vinogradov, Alexander Boukhanovsky, Nikolay Nikitin, Andrei Dmitrenko, Anna Kalyuzhnaya and Andrey Savchenko	6956
<i>PersonaGym: Evaluating Persona Agents and LLMs</i>	
Vinay Samuel, Henry Peng Zou, Yue Zhou, Shreyas Chaudhari, Ashwin Kalyan, Tanmay Rajpurohit, Ameet Deshpande, Karthik R Narasimhan and Vishvak Murahari	6999
<i>LM2Protein: A Structure-to-Token Protein Large Language Model</i>	
Chang Zhou, Yuheng Shan, Pengan Chen, Xiangyu Shi, Zikang Wang, Yanting Li and Jiyue Jiang	7023
<i>How Well Can Reasoning Models Identify and Recover from Unhelpful Thoughts?</i>	
Sohee Yang, Sang-Woo Lee, Nora Kassner, Daniela Gottesman, Sebastian Riedel and Mor Geva	7030
<i>From Token to Action: State Machine Reasoning to Mitigate Overthinking in Information Retrieval</i>	
Dohyeon Lee, Yeonseok Jeong and Seung-won Hwang	7048
<i>Locate-then-Merge: Neuron-Level Parameter Fusion for Mitigating Catastrophic Forgetting in Multi-modal LLMs</i>	
Zeping Yu and Sophia Ananiadou	7065
<i>Improving Influence-based Instruction Tuning Data Selection for Balanced Learning of Diverse Capabilities</i>	
Qirun Dai, Dylan Zhang, Jiaqi W. Ma and Hao Peng	7079
<i>Diagnosing Moral Reasoning Acquisition in Language Models: Pragmatics and Generalization</i>	
Guangliang Liu, Zimo Qi, Xitong Zhang, Lei Jiang and Kristen Johnson	7103
<i>Discourse Heuristics For Paradoxically Moral Self-Correction</i>	
Guangliang Liu, Zimo Qi, Xitong Zhang and Kristen Johnson	7118

<i>Invisible Prompts, Visible Threats: Malicious Font Injection in External Resources for Large Language Models</i>	
Junjie Xiong, Changjia Zhu, Shuhang Lin, Chong Zhang, Yongfeng Zhang, Yao Liu and Lingyao Li	7133
<i>Turning the Tide: Repository-based Code Reflection</i>	
Wei Zhang, Jian Yang, Jiayi Yang, Ya Wang, Zhoujun Li, Zeyu Cui, Binyuan Hui and Junyang Lin	7148
<i>Reinforcement Learning with Supervised Alignment</i>	
João Luís Lins and Jia Xu	7165
<i>EmByte: Decomposition and Compression Learning for Small yet Private NLP</i>	
Shenglan Li, Jia Xu and Mengjiao Zhang	7182
<i>GUARD: Glocal Uncertainty-Aware Robust Decoding for Effective and Efficient Open-Ended Text Generation</i>	
Yuanhao Ding, Esteban Garces Arias, Meimingwei Li, Julian Rodemann, Matthias Aßenmacher, Danlu Chen, Gaojuan Fan, Christian Heumann and Chongsheng Zhang	7202
<i>Efficiently Editing Mixture-of-Experts Models with Compressed Experts</i>	
Yifei He, Yang Liu, Chen Liang and Hany Hassan Awadalla	7227
<i>FinGEAR: Financial Mapping-Guided Enhanced Answer Retrieval</i>	
Ying Li, Mengyu Wang, Miguel de Carvalho, Sotirios Sabanis and Tiejun Ma	7239
<i>FM2DS: Few-Shot Multimodal Multihop Data Synthesis with Knowledge Distillation for Question Answering</i>	
Amirhossein Abaskohi, Spandana Gella, Giuseppe Carenini and Issam H. Laradji	7256
<i>SQUARE: Unsupervised Retrieval Adaptation via Synthetic Data</i>	
Jinsung Yoon, Junhao Zeng and Sercan O Arik	7283
<i>Knowledge-enhanced Multimodal ECG Representation Learning with Arbitrary-Lead Inputs</i>	
Che Liu, Cheng Ouyang, Zhongwei Wan, Haozhe Wang, Wenjia Bai and Rossella Arcucci ..	7298
<i>Seeing Race, Feeling Bias: Emotion Stereotyping in Multimodal Language Models</i>	
Mahammed Kamruzzaman, Amanda Cercas Curry, Alba Cercas Curry and Flor Miriam Plaza-del-Arco	7317
<i>AdaptMerge: Inference Time Adaptive Visual and Language-Guided Token Merging for Efficient Large Multimodal Models</i>	
Zahidul Islam and Mrigank Rochan	7352
<i>Federated Retrieval-Augmented Generation: A Systematic Mapping Study</i>	
Abhijit Chakraborty, Chahana Dahal and Vivek Gupta	7362
<i>A Survey of Pun Generation: Datasets, Evaluations and Methodologies</i>	
Yuchen Su, Yonghua Zhu, Ruofan Wang, Zijian Huang, Diana Benavides-Prado and Michael J. Witbrock	7375
<i>Evaluating the Robustness and Accuracy of Text Watermarking Under Real-World Cross-Lingual Manipulations</i>	
Mansour Al Ghanim, Jiaqi Xue, Rochana Prih Hastuti, Mengxin Zheng, Yan Solihin and Qian Lou	7396

<i>HDiff: Confidence-Guided Denoising Diffusion for Robust Hyper-relational Link Prediction</i> Xiangfeng Luo, Ruoxin Zheng, Jianqiang Huang and Hang Yu	7417
<i>Spotlighter: Revisiting Prompt Tuning from a Representative Mining View</i> Yutong Gao, Maoyuan Shao, Xinyang Huang, Chuang Zhu, Yu Weng, Xuan Liu, Lijuan Sun and Guoshun Nan	7435
<i>Offloaded Reasoning: Efficient Inference for Large Language Models via Modular Reasoning and Refinement</i> Ishan Jindal, Jayant Taneja, Badrinath Chandana, Vikas Kapur and Sachin Dev Sharma	7450
<i>Wait, We Don't Need to Wait"! Removing Thinking Tokens Improves Reasoning Efficiency</i> Chenlong Wang, Yuanning Feng, Dongping Chen, Zhaoyang Chu, Ranjay Krishna and Tianyi Zhou	7459
<i>Towards Reverse Engineering of Language Models: A Survey</i> Xinpeng Ti, Wentao Ye, Zhifang Zhang, Junbo Zhao, Chang Yao, Lei Feng and Haobo Wang	7483
<i>LIFTED: Multimodal Clinical Trial Outcome Prediction via Large Language Models and Mixture-of-Experts</i> Wenhao Zheng, Liaoyaqi Wang, Dongshen Peng, Hongxia Xu, Yun Li, Hongtu Zhu, Tianfan Fu and Huaxiu Yao	7503
<i>Addition in Four Movements: Mapping Layer-wise Information Trajectories in LLMs</i> Yao Yan	7518
<i>CoMoE: Contrastive Representation for Mixture-of-Experts in Parameter-Efficient Fine-tuning</i> Jinyuan Feng, ChaoPeng Wei, Tenghai Qiu, Tianyi Hu and Zhiqiang Pu	7533
<i>GuiLoMo: Allocating Experts and Ranks for LoRA-MoE via Bilevel Optimization with Guided Selection Vectors</i> Xinrong Chen, Hengyuan Zhang, Yingmin Qiu, Xiao Liang, Ziyue Li, Guanyu Wang, Weiping Li, Tong Mo, Hayden Kwok-Hay So and Ngai Wong	7552
<i>Rotate, Clip, and Partition: Towards W2A4KV4 Quantization by Integrating Rotation and Learnable Non-uniform Quantizer</i> Euntae Choi, Sumin Song, Woosang Lim and Sungjoo Yoo	7568
<i>Decoding in Latent Spaces for Efficient Inference in LLM-based Recommendation</i> Chengbing Wang, Yang Zhang, Zhicheng Wang, Tianhao Shi, Keqin Bao, Fuli Feng and Tat-Seng Chua	7591
<i>Forget for Get: A Lightweight Two-phase Gradient Method for Knowledge Editing in Large Language Models</i> Yanhong Li, Min Yang, Xiping Hu and Chengming Li	7604
<i>AutoEvolve: Automatically Evolving Queries for Applicable and Scalable Retrieval-Augmented Generation Benchmarking</i> Ding-Chu Zhang, Xiaowen Zhang, Yue Fei, Renjun Hu, Xiao-Wen Yang, Zhi Zhou, Baixuan Li, Yu-Feng Li, Xing Shi and Wei Lin	7624
<i>Temporal Alignment of Time Sensitive Facts with Activation Engineering</i> Sanjay Govindan, Maurice Pagnucco and Yang Song	7640
<i>ChronoBias: A Benchmark for Evaluating Temporal Group Bias in the Time-sensitive Knowledge of Large Language Models</i> Kyungmin Kim, Youngbin Choi, Hyounghun Kim, Dongwoo Kim and Sangdon Park	7658

<i>MC²: A Minimum-Coverage and Dataset-Agnostic Framework for Compositional Generalization of LLMs on Semantic Parsing</i>	
Ziyao Xu, Zhe Yang and Houfeng Wang	7694
<i>Learning to Instruct: Fine-Tuning a Task-Aware Instruction Optimizer for Black-Box LLMs</i>	
Yunzhe Qi, Jinjin Tian, Tianci Liu, Ruirui Li, Tianxin Wei, Hui Liu, Xianfeng Tang, Monica Xiao Cheng and Jingrui He	7707
<i>Enriching Patent Claim Generation with European Patent Dataset</i>	
Lekang Jiang, Chengzu Li and Stefan Goetz	7734
<i>StepKE: Stepwise Knowledge Editing for Multi-Hop Question Answering</i>	
Jaewook Lee, Dahyun Jung and Heuseok Lim	7752
<i>AutoDCWorkflow: LLM-based Data Cleaning Workflow Auto-Generation and Benchmark</i>	
Lan Li, Liri Fang, Bertram Ludäscher and Vetle I Torvik	7766
<i>Hidden Ghost Hand: Unveiling Backdoor Vulnerabilities in MLLM-Powered Mobile GUI Agents</i>	
Pengzhou Cheng, Haowen Hu, Zheng Wu, Zongru Wu, Tianjie Ju, Daizong Ding, Zhuosheng Zhang and Gongshen Liu	7781
<i>Scale Down to Speed Up: Dynamic Data Selection for Reinforcement Learning</i>	
Zhuoyue Chen, Jihai Zhang, Ben Liu, Fangquan Lin and Wotao Yin	7806
<i>Towards Efficient CoT Distillation: Self-Guided Rationale Selector for Better Performance with Fewer Rationales</i>	
Jianzhi Yan, Le Liu, Youcheng Pan, Shiwei Chen, Yang Xiang and Buzhou Tang	7818
<i>GeoDANO: Geometric VLM with Domain Agnostic Vision Encoder</i>	
Seunghyuk Cho, Zhenyue Qin, Yang Liu, Youngbin Choi, Seungbeom Lee and Dongwoo Kim	7836
<i>Leveraging 3D Gaussian for Temporal Knowledge Graph Embedding</i>	
Jiang Li, Xiangdong Su and Guanglai Gao	7852
<i>LLMAP: LLM-Assisted Multi-Objective Route Planning with User Preferences</i>	
Liangqi Yuan, Dong-Jun Han, Christopher Brinton and Sabine Brunswicker	7866
<i>ZEBRA: Leveraging Model-Behavioral Knowledge for Zero-Annotation Preference Dataset Construction</i>	
Jeesu Jung, Chanjun Park and Sangkeun Jung	7895
<i>Token Knowledge: A New Perspective For Knowledge in Large Language Models</i>	
Jieyong Wang, Chunyao Song and Tingjian Ge	7912
<i>Adaptive Schema-aware Event Extraction with Retrieval-Augmented Generation</i>	
Sheng Liang, Hang Lv, Zhihao Wen, Yaxiong Wu, Yongyue Zhang, Hao Wang and Yong Liu	7927
<i>Enhancing Attributed Question Answering using Tailored Progressive Curriculum Learning</i>	
Yuhan Chen, Bawei Zou, Yifan Fan, Yuchong Chen, Shujun Cao and Yu Hong	7947
<i>REAR: Reinforced Reasoning Optimization for Event Argument Extraction with Relation-Aware Support</i>	
Jianwen Luo, Yu Hong, Shuai Yang and Jianmin Yao	7957
<i>COMI-LINGUA: Expert Annotated Large-Scale Dataset for Multitask NLP in Hindi-English Code-Mixing</i>	
Rajvee Sheth, Himanshu Beniwal and Mayank Singh	7973

<i>Nine Ways to Break Copyright Law and Why Our LLM Won't: A Fair Use Aligned Generation Framework</i>	
Aakash Sen Sharma, Debdeep Sanyal, Priyansh Srivastava, Sundar Athreya H, Shirish Karande, Mohan Kankanhalli and Murari Mandal	7993
<i>InteractSpeech: A Speech Dialogue Interaction Corpus for Spoken Dialogue Model</i>	
Yifu Chen, Shengpeng Ji, Ziqing Wang, Hanting Wang and Zhou Zhao	8024
<i>Enhancing SQL Table Acquisition with Reverse Engineering for Text-to-SQL</i>	
Shixin Liu, Haoyu Xu and Yu Hong	8034
<i>DynamicKV: Task-Aware Adaptive KV Cache Compression for Long Context LLMs</i>	
Xiabin Zhou, Wenbin Wang, Minyan Zeng, Jiaxian Guo, Xuebo Liu, Li Shen, Min Zhang and Liang Ding	8042
<i>ASD-iLLM: An Intervention Large Language Model for Autistic Children based on Real Clinical Dialogue Intervention Dataset</i>	
Shuzhong Lai, Chenxi Li, Junhong Lai, Yucun Zhong, Chenyu Yan, Xiang Li, Haifeng Li, Gang Pan, Lin Yao and Yueming Wang	8058
<i>GDLLM: A Global Distance-aware Modeling Approach Based on Large Language Models for Event Temporal Relation Extraction</i>	
Jie Zhao, Wanting Ning, Yuxiao Fei, Yubo Feng and Lishuang Li	8080
<i>More Tokens, Lower Precision: Towards the Optimal Token-Precision Trade-off in KV Cache Compression</i>	
Jiebin Zhang, Dawei Zhu, Yifan Song, Wenhao Wu, Chuqiao Kuang, Xiaoguang Li, Lifeng Shang, Qun Liu and Sujian Li	8092
<i>cAST: Enhancing Code Retrieval-Augmented Generation with Structural Chunking via Abstract Syntax Tree</i>	
Yilin Zhang, Xinran Zhao, Zora Zhiruo Wang, Chenyang Yang, Jiayi Wei and Tongshuang Wu	8106
<i>A Group Fairness Lens for Large Language Models</i>	
Guanqun Bi, Yuqiang Xie, Lei Shen and Yanan Cao	8117
<i>VLM Is a Strong Reranker: Advancing Multimodal Retrieval-augmented Generation via Knowledge-enhanced Reranking and Noise-injected Training</i>	
Zhanpeng Chen, Chengjin Xu, Yiyang Qi, Xuhui Jiang and Jian Guo	8140
<i>Rethinking DPO: The Role of Rejected Responses in Preference Misalignment</i>	
Jae Hyeon Cho, JunHyeok Oh, Myunsoo Kim and Byung-Jun Lee	8159
<i>Enhancing Recommendation Explanations through User-Centric Refinement</i>	
Jingsen Zhang, Zihang Tian, Xueyang Feng, Xu Chen and Chong Chen	8177
<i>Distributional Surgery for Language Model Activations</i>	
Bao Nguyen, Binh Nguyen, Duy Nguyen and Viet Anh Nguyen	8192
<i>Improving Alignment in LVLMs with Debiased Self-Judgment</i>	
Sihan Yang, Chenhang Cui, Zihao Zhao, Yiyang Zhou, Weilong Yan, Ying Wei and Huaxiu Yao	8213
<i>Low-Confidence Gold: Refining Low-Confidence Samples for Efficient Instruction Tuning</i>	
Hongyi Cai, Jie li, Mohammad Mahdinur Rahman and Wenzhen Dong	8233

<i>Safeguarding Privacy of Retrieval Data against Membership Inference Attacks: Is This Query Too Close to Home?</i>	
Yujin Choi, Youngjoo Park, Junyoung Byun, Jaewook Lee and Jinseong Park	8241
<i>Causal-LLM: A Unified One-Shot Framework for Prompt- and Data-Driven Causal Graph Discovery</i>	
Amartya Roy, N Devharish, Shreya Ganguly and Kripabandhu Ghosh	8259
<i>LRPLAN: A Multi-Agent Collaboration of Large Language and Reasoning Models for Planning with Implicit & Explicit Constraints</i>	
T Karthikeyan, Om Dehlan, Mausam and Manish Gupta	8280
<i>DLPO: Towards a Robust, Efficient, and Generalizable Prompt Optimization Framework from a Deep-Learning Perspective</i>	
Dengyun Peng, Yuhang Zhou, Qiguang Chen, JinHao Liu, Jingjing Chen, Libo Qin and Wanxiang Che	8311
<i>Towards Robust Few-Shot Relation Classification: Incorporating Relation Description with Agreement</i>	
Mengting Hu, Jianfeng Wu, Ming Jiang, Yalan Xie, Zhunheng Wang, Rui Ying, Xiaoyi Liu, Ruixuan Xu, Hang Gao and Renhong Cheng	8335
<i>For a Fistful of Puns: Evaluating a Puns in Multiword Expressions Identification Algorithm Without Dedicated Dataset</i>	
Julien Bezançon and Gaël Lejeune	8350
<i>Watermarking for Factuality: Guiding Vision-Language Models Toward Truth via Tri-layer Contrastive Decoding</i>	
Kyungryul Back, Seongbeom Park, Milim Kim, Mincheol Kwon, SangHyeok Lee, Hyunyoung Lee, Junhee Cho, Seunghyun Park and Jinkyu Kim	8371
<i>Are the Reasoning Models Good at Automated Essay Scoring?</i>	
Lui Yoshida	8388
<i>Rethinking LLM-Based Recommendations: A Personalized Query-Driven Parallel Integration</i>	
Donghee Han, Hwanjun Song and Mun Yong Yi	8395
<i>RefVNL: Towards Scalable Evaluation of Subject-driven Text-to-image Generation</i>	
Aviv Slobodkin, Hagai Taitelbaum, Yonatan Bitton, Brian Gordon, Michal Sokolik, Nitzan Bitton Guetta, Almog Gueta, Royi Rassin, Dani Lischinski and Idan Szpektor	8420
<i>What data should I include in my POS tagging training set?</i>	
Zoey Liu, Masoud Jasbi, Christan Grant, Kenji Sagae and Emily Prud’hommeaux	8439
<i>AttnComp: Attention-Guided Adaptive Context Compression for Retrieval-Augmented Generation</i>	
Lvzhou Luo, Yixuan Cao and Ping Luo	8456
<i>SafeInt: Shielding Large Language Models from Jailbreak Attacks via Safety-Aware Representation Intervention</i>	
Jiaqi Wu, Chen Chen, Chunyan Hou and Xiaojie Yuan	8473
<i>Staged Knowledge Distillation Through Least-to-Most Prompting: Optimizing Teacher Guidance via Difficulty-Aware Training</i>	
Mengxiang Zhang and Lingyuan Liu	8489
<i>LLM Distillation for Efficient Few-Shot Multiple Choice Question Answering</i>	
Patrick Sutanto, Joan Santoso, Esther Irawati Setiawan and Aji Prasetya Wibawa	8502

<i>Teaching LLMs to Plan, Not Just Solve: Plan Learning Boosts LLMs Generalization in Reasoning Tasks</i> Tianlong Wang, Junzhe Chen, Weibin Liao, Xueting Han and Jing Bai	8531
<i>FedCoT: Federated Chain-of-Thought Distillation for Large Language Models</i> Tao Fan, Weijing Chen, Yan Kang, Guoqiang Ma, Hanlin Gu, Yuanfeng Song, Lixin Fan and Qiang Yang	8546
<i>SalaMAnder: Shapley-based Mathematical Expression Attribution and Metric for Chain-of-Thought Reasoning</i> Yue Xin, Chen Shen, Shaotian Yan, Xiaosong Yuan, Yaoming Wang, Xiaofeng Zhang, Chenxi Huang and Jieping Ye	8558
<i>Representing LLMs in Prompt Semantic Task Space</i> Idan Kashani, Avi Mendelson and Yaniv Nemcovsky	8578
<i>PersLLM: A Personified Training Approach for Large Language Models</i> Zheni Zeng, Jiayi Chen, Huimin Chen, Yukun Yan, Yuxuan Chen, Zhenghao Liu, Zhiyuan Liu and Maosong Sun	8598
<i>The Illusion of Randomness: How LLMs Fail to Emulate Stochastic Decision-Making in Rock-Paper-Scissors Games?</i> Zihao Guo, Hongtao Lv, Chaoli Zhang, Yibowen Zhao, Yixin Zhang and Lizhen Cui	8618
<i>DAPE-BR: Distance-Aware Positional Encoding for Mitigating Object Hallucination in LVLMS</i> Mingrui Xie, Tianxiang Xu, Qianhai Tang, Shanming Yao, Xiaofeng Zhang and Junliang Du	8638
<i>From Confidence to Collapse in LLM Factual Robustness</i> Alina Fastowski, Bardh Prenkaj and Gjergji Kasneci	8650
<i>CtrlNews: LLM-based Multi-Agent Controllable News Writing via Knowledge Gravitational Field</i> Yifei Xu, Yingjie Zong, Wang Zhonghua, Sirui Wu, Yuan Rao, Dan Zhang and Shuiguang Deng	8668
<i>Joint Enhancement of Relational Reasoning for Long-Context LLMs</i> Zhirui Chen, Wei Shen, Jiashui Huang and Ling Shao	8706
<i>Training Medical QA Models Based on Mixed Rewards from Multiple-Choice and Open-Ended Questions</i> Yue Qiu, Yujan Ting, Pei Dong, Terrence Chen and Weijing Huang	8721
<i>Rethink Rumor Detection in the Era of LLMs: A Review</i> Chang Yang, Peng Zhang, Jing Zhang, Hui Gao and Changhao Song	8730
<i>ScholarBench: A Bilingual Benchmark for Abstraction, Comprehension, and Reasoning Evaluation in Academic Contexts</i> Dongwon Noh, Donghyeok Koh, Junghun Yuk, Gyuwan Kim, Jae Yong Lee, KyungTae Lim and Cheoneum Park	8750
<i>MAGIC: A Multi-Hop and Graph-Based Benchmark for Inter-Context Conflicts in Retrieval-Augmented Generation</i> Jungyeon Lee, Lee Kangmin and Taek Kim	8783
<i>Align Attention Heads Before Merging Them: An Effective Way for Converting MHA to GQA</i> Qingyun Jin, Xiaohui Song, Feng Zhou and Zengchang Qin	8804

<i>DRBO: Mitigating Short Board Effect via Dynamic Reward Balancing in Multi-reward LLM Optimization</i>	
Nuo Chen, Yufei Gao, Yongnan Jin, Yan Hu, Anningzhe Gao, Lingyong Yan and Benyou Wang	8817
<i>Enhancing LLM Knowledge Learning through Generalization</i>	
Mingkang Zhu, Xi Chen, Zhongdao Wang, Bei Yu, Hengshuang Zhao and Jiaya Jia	8842
<i>FastCuRL: Curriculum Reinforcement Learning with Stage-wise Context Scaling for Efficient Training RI-like Reasoning Models</i>	
Mingyang Song, Mao Zheng, Zheng Li, Wenjie Yang and Xuan Luo	8856
<i>TR-MTEB: A Comprehensive Benchmark and Embedding Model Suite for Turkish Sentence Representations</i>	
Mehmet Selman Baysan and Tunga Gungor	8867
<i>ImpRAG: Retrieval-Augmented Generation with Implicit Queries</i>	
Wenzheng Zhang, Xi Victoria Lin, Karl Stratos, Wen-tau Yih and Mingda Chen	8888
<i>HEAL: A Hypothesis-Based Preference-Aware Analysis Framework</i>	
Yifu Huo, Chenglong Wang, Qiren Zhu, Shunjie Xing, Tong Xiao, Chunliang Zhang, Tongran Liu and JingBo Zhu	8901
<i>A Survey of Multilingual Reasoning in Language Models</i>	
Akash Ghosh, Debayan Datta, Sriparna Saha and Chirag Agarwal	8920
<i>CLEAR: A Framework Enabling Large Language Models to Discern Confusing Legal Paragraphs</i>	
Qi Xu, Qian Liu, Hao Fei, Hang Yu, Shuhao Guan and Xiao Wei	8937
<i>NAP²: A Benchmark for Naturalness and Privacy-Preserving Text Rewriting by Learning from Human</i>	
Shuo Huang, William Maclean, Xiaoxi Kang, Qionikai Xu, Zhuang Li, Xingliang Yuan, Gholamreza Haffari and Lizhen Qu	8954
<i>Chain of Ideas: Revolutionizing Research Via Novel Idea Development with LLM Agents</i>	
Long Li, Weiwen Xu, Jiayan Guo, Ruo Chen Zhao, Xingxuan Li, Yuqian Yuan, Boqiang Zhang, Yuming Jiang, Yifei Xin, Ronghao Dang, Yu Rong, Deli Zhao, Tian Feng and Lidong Bing	8971
<i>Unveiling Multimodal Processing: Exploring Activation Patterns in Multimodal LLMs for Interpretability and Efficiency</i>	
Chuan Wu, Meng Su, Youxuan Fang and Shaolin Zhu	9005
<i>Self-Supervised Prompt Optimization</i>	
Jinyu Xiang, Jiayi Zhang, Zhaoyang Yu, Xinbing Liang, Fengwei Teng, Jinhao Tu, Fashen Ren, Xiangru Tang, Sirui Hong, Chenglin Wu and Yuyu Luo	9017
<i>Polish-English medical knowledge transfer: A new benchmark and results</i>	
Łukasz Grzybowski, Jakub Pokrywka, Michał Ciesiółka, Jeremi Ignacy Kaczmarek and Marek Kubis	9042
<i>Hard Negatives, Hard Lessons: Revisiting Training Data Quality for Robust Information Retrieval with LLMs</i>	
Nandan Thakur, Crystina Zhang, Xueguang Ma and Jimmy Lin	9064
<i>EventRelBench: A Comprehensive Benchmark for Evaluating Event Relation Understanding in Large Language Models</i>	
Jie Gong, Biaoshuai Zheng and Qiwang hu	9084

<i>S2LPP: Small-to-Large Prompt Prediction across LLMs</i>	
Liang Cheng, Tianyi Li, Zhaowei Wang and Mark Steedman	9100
<i>DroidCall: A Dataset for LLM-powered Android Intent Invocation</i>	
Weikai Xie, Li Zhang, Shihe Wang, Rongjie Yi and Mengwei Xu	9116
<i>Tool Zero: Training Tool-Augmented LLMs via Pure RL from Scratch</i>	
Yirong Zeng, Xiao Ding, Yutai Hou, Yuxian Wang, Li Du, Juyi Dai, Qiuyang Ding, Duyu Tang, Dandan Tu, Weiwen Liu, Bing Qin and Ting Liu	9135
<i>INREACT: An Inspire-Then-Reinforce Training Framework For Multimodal GUI Agent</i>	
Yuanlei Wang, Liuzhou Zhang, Haohao Luo and Ying Shen	9148
<i>Facts Fade Fast: Evaluating Memorization of Outdated Medical Knowledge in Large Language Models</i>	
Juraj Vladika, Mahdi Dhaini and Florian Matthes	9161
<i>Zero-Shot Privacy-Aware Text Rewriting via Iterative Tree Search</i>	
Shuo Huang, Xingliang Yuan, Gholamreza Haffari and Lizhen Qu	9175
<i>KoLEG: On-the-Fly Korean Legal Knowledge Editing with Continuous Retrieval</i>	
Jaehyung Seo, Dahyun Jung, Jaewook Lee, Yongchan Chun, Dongjun Kim, Hwijung Ryu, Donghoon Shin and Heuseok Lim	9191
<i>HARE: an entity and relation centric evaluation framework for histopathology reports</i>	
Yunsoo Kim, Michal Wen Sheue Ong, Alex Shavick, Honghan Wu and Adam P. Levine ...	9218
<i>VeriFastScore: Speeding up long-form factuality evaluation</i>	
Rishanth Rajendhran, Amir Zadeh, Matthew Sarte, Chuan Li and Mohit Iyyer	9234
<i>B-REASO: A Multi-Level Multi-Faceted Bengali Evaluation Suite for Foundation Models</i>	
Md Tanzib Hosain and Md Kishor Morol	9260
<i>Extracting Conceptual Spaces from LLMs Using Prototype Embeddings</i>	
Nitesh Kumar, Usashi Chatterjee and Steven Schockaert	9275
<i>FC-Attack: Jailbreaking Multimodal Large Language Models via Auto-Generated Flowcharts</i>	
Ziyi Zhang, Zhen Sun, Zongmin Zhang, Jihui Guo and Xinlei He	9299
<i>Multilingual Data Filtering using Synthetic Data from Large Language Models</i>	
Jonas Waldendorf, Barry Haddow, Alexandra Birch and Mateusz Klimaszewski	9317
<i>SAFE: A Sparse Autoencoder-Based Framework for Robust Query Enrichment and Hallucination Mitigation in LLMs</i>	
Samir Abdaljalil, Filippo Pallucchini, Andrea Seveso, Hasan Kurban, Fabio Mercorio and Erchin Serpedin	9335
<i>Soteria: Language-Specific Functional Parameter Steering for Multilingual Safety Alignment</i>	
Somnath Banerjee, Sayan Layek, Pratyush Chatterjee, Animesh Mukherjee and Rima Hazra	9347
<i>LLMs as a synthesis between symbolic and distributed approaches to language</i>	
Gemma Boleda	9365
<i>MIND: Towards Immersive Psychological Healing with Multi-Agent Inner Dialogue</i>	
Yujia Chen, Changsong Li, Yiming Wang, Tianjie Ju, Qingqing Xiao, Nan Zhang, Zifan Kong, Peng Wang and Binyu Yan	9380

<i>A Monte-Carlo Sampling Framework For Reliable Evaluation of Large Language Models Using Behavioral Analysis</i>	
Davood Wadi and Marc Fredette	9414
<i>Understanding How Value Neurons Shape the Generation of Specified Values in LLMs</i>	
Yi Su, Jiayi Zhang, Shu Yang, Xinhai Wang, Lijie Hu and Di Wang	9433
<i>Likelihood Variance as Text Importance for Resampling Texts to Map Language Models</i>	
Momose Oyama, Ryo Kishino, Hiroaki Yamagiwa and Hidetoshi Shimodaira	9453
<i>Think Twice, Generate Once: Safeguarding by Progressive Self-Reflection</i>	
Hoang Phan, Victor Li and Qi Lei	9466
<i>Efficient Integration of External Knowledge to LLM-based World Models via Retrieval-Augmented Generation and Reinforcement Learning</i>	
Chang Yang, Xinrun Wang, Qinggang Zhang, Qi Jiang and Xiao Huang	9484
<i>Comparing Apples to Oranges: A Dataset & Analysis of LLM Humour Understanding from Traditional Puns to Topical Jokes</i>	
Tyler Loakman, William Thorne and Chenghua Lin	9502
<i>Modeling, Evaluating, and Embodying Personality in LLMs: A Survey</i>	
Iago Alves Brito, Julia Soares Dollis, Fernanda Bufon Färber, Pedro Schindler Freire Brasil Ribeiro, Rafael Teixeira Sousa and Arlindo Rodrigues Galvão Filho	9519
<i>Benchmarking the Detection of LLMs-Generated Modern Chinese Poetry</i>	
Shanshan Wang, Junchao Wu, Fengying Ye, Derek F. Wong, Jingming Yao and Lidia S. Chao	9533
<i>Leveraging the Cross-Domain & Cross-Linguistic Corpus for Low Resource NMT: A Case Study On Bhili-Hindi-English Parallel Corpus</i>	
Pooja Singh, Shashwat Bhardwaj, Vaibhav Sharma and Sandeep Kumar	9553
<i>Creative Preference Optimization</i>	
Mete Ismayilzada, Antonio Laverghetta Jr., Simone A. Luchini, Reet Patel, Antoine Bosselut, Lonke Van Der Plas and Roger E. Beaty	9580
<i>Assistant-Guided Mitigation of Teacher Preference Bias in LLM-as-a-Judge</i>	
Zhuo Liu, Moxin Li, Xun Deng, Qifan Wang and Fuli Feng	9610
<i>Uplift-RAG: Uplift-Driven Knowledge Preference Alignment for Retrieval-Augmented Generation</i>	
Changle Qu, Sunhao Dai, Hengyi Cai, Yiyang Cheng, Jun Xu, Shuaiqiang Wang and Dawei Yin	9632
<i>Sugar-Coated Poison: Benign Generation Unlocks Jailbreaking</i>	
Yuhang Wu, Yu-Jie Xiong, Hao Zhang, Jia-Chen Zhang and Zheng Zhou	9645
<i>DivScene: Towards Open-Vocabulary Object Navigation with Large Vision Language Models in Diverse Scenes</i>	
Zhaowei Wang, Hongming Zhang, Tianqing Fang, Ye Tian, Yue Yang, Kaixin Ma, Xiaoman Pan, Yangqiu Song and Dong Yu	9666
<i>Data-scarce Behavior Editing of Language Models</i>	
Joykirat Singh, Subhabrata Dutta and Tanmoy Chakraborty	9687
<i>FIER: Fine-Grained and Efficient KV Cache Retrieval for Long-context LLM Inference</i>	
Dongwei Wang, Zijie Liu, Song Wang, Yuxin Ren, Jianing Deng, Jingtong Hu, Tianlong Chen and Huanrui Yang	9702

<i>SVeritas: Benchmark for Robust Speaker Verification under Diverse Conditions</i>	
Massa Baali, Sarthak Bisht, Francisco Teixeira, Kateryna Shapovalenko, Rita Singh and Bhiksha Raj	9714
<i>CAARMA: Class Augmentation with Adversarial Mixup Regularization</i>	
Massa Baali, Xiang Li, Hao Chen, Syed Abdul Hannan, Rita Singh and Bhiksha Raj	9732
<i>Bringing Pedagogy into Focus: Evaluating Virtual Teaching Assistants' Question-Answering in Asynchronous Learning Environments</i>	
Li Siyan, Zhen Xu, Vethavikashini Chithrara Raghuram, Xuanming Zhang, Renzhe Yu and Zhou Yu	9743
<i>Demystifying Multilingual Reasoning in Process Reward Modeling</i>	
Weixuan Wang, Minghao Wu, Barry Haddow and Alexandra Birch	9775
<i>BehaviorSFT: Behavioral Token Conditioning for Health Agents Across the Proactivity Spectrum</i>	
Yubin Kim, Zhiyuan Hu, Hyewon Jeong, Eugene W Park, Shuyue Stella Li, Chanwoo Park, Shiyun Xiong, MingYu Lu, Hyeonhoon Lee, Xin Liu, Daniel McDuff, Cynthia Breazeal, Samir Tulebaev and Hae Won Park	9789
<i>LaMP-Cap: Personalized Figure Caption Generation With Multimodal Figure Profiles</i>	
Ho Yin Sam Ng, Edward Hsu, Aashish Anantha Ramakrishnan, Branislav Kveton, Nedim Lipka, Franck Dernoncourt, Dongwon Lee, Tong Yu, Sungchul Kim, Ryan A. Rossi and Ting-Hao Kenneth Huang	9818
<i>Efficient Dynamic Clustering-Based Document Compression for Retrieval-Augmented-Generation</i>	
Weitao Li, Xiangyu Zhang, Kaiming Liu, Xuanyu Lei, Weizhi Ma and Yang Liu	9833
<i>HebID: Detecting Social Identities in Hebrew-language Political Text</i>	
Guy Mor-Lan, Naama Rivlin-Angert, Yael R. Kaplan, Tamir Sheaffer and Shaul R. Shenhav	9850
<i>Dub-S2ST: Textless Speech-to-Speech Translation for Seamless Dubbing</i>	
Jeongsoo Choi, Jaehun Kim and Joon Son Chung	9871
<i>FinGrAct: A Framework for FINE-GRained Evaluation of ACTIONability in Explainable Automatic Fact-Checking</i>	
Islam Eldifrawi, Shengrui Wang and Amine Trabelsi	9882
<i>What Has Been Lost with Synthetic Evaluation?</i>	
Alexander Gill, Abhilasha Ravichander and Ana Marasovic	9902
<i>Bold Claims or Self-Doubt? Factuality Hallucination Type Detection via Belief State</i>	
Dongyu Zhang, Qingqing Hong, Bingxuan Hou, Jiayi Lin, Chenyang Zhang, Jialin Li and Junli Wang	9946
<i>Proxy Barrier: A Hidden Repeater Layer Defense Against System Prompt Leakage and Jailbreaking</i>	
Pedro Schindler Freire Brasil Ribeiro, Iago Alves Brito, Rafael Teixeira Sousa, Fernanda Bufon Färber, Julia Soares Dollis and Arlindo Rodrigues Galvão Filho	9960
<i>AraSafe: Benchmarking Safety in Arabic LLMs</i>	
Hamdy Mubarak, Abubakr Mohamed and Majd Hawasly	9976
<i>Nested Named Entity Recognition as Single-Pass Sequence Labeling</i>	
Alberto Muñoz-Ortiz, David Vilares, Caio Corro and Carlos Gómez-Rodríguez	9993
<i>DeCoRe: Decoding by Contrasting Retrieval Heads to Mitigate Hallucinations</i>	
Aryo Pradipta Gema, Chen Jin, Ahmed Abdulaal, Tom Diethe, Philip Alexander Teare, Beatrice Alex, Pasquale Minervini and Amrutha Saseendran	10003

<i>Catch Me If You Can? Not Yet: LLMs Still Struggle to Imitate the Implicit Writing Styles of Everyday Authors</i>	
Zhengxiang Wang, Nafis Irtiza Tripto, Solha Park, Zhenzhen Li and Jiawei Zhou	10040
<i>Fine-Tuning Encoder-Decoder Models with Contrastive Learning for In-Context Distractor Generation</i>	
Elaf Alhazmi, Quan Z. Sheng, Wei Emma Zhang, Mohammed I. Thanoon, Haojie Zhuang, Behnaz Soltani and Munazza Zaib	10056
<i>Conflicts in Texts: Data, Implications and Challenges</i>	
Siyi Liu and Dan Roth	10073
<i>Recognizing Limits: Investigating Infeasibility in Large Language Models</i>	
Wenbo Zhang, Zihang Xu and Hengrui Cai	10092
<i>VQA-Augmented Machine Translation with Cross-Modal Contrastive Learning</i>	
Zhihui Zhang, Shiliang Sun, Jing Zhao, Tengfei Song and Hao Yang	10113
<i>Learning to Describe Implicit Changes: Noise-robust Pre-training for Image Difference Captioning</i>	
Zixin Guo, Jiayang Sun, Tzu-Jui Julius Wang, Abduljalil Radman, Selen Pehlivan, Min Cao and Jorma Laaksonen	10125
<i>SOLAR: Serendipity Optimized Language Model Aligned for Recommendation</i>	
Zichen Yuan, Lifan Sun, Yucen Zhuang, Yue Wang, Xinyuan Song, Tianqi Xu, Siyuan Li, Junchen Fu, Youhua Li, Sirui Hong, Jiaqi Chen, Joemon M. Jose and Yongxin Ni	10146
<i>AIRepr: An Analyst-Inspector Framework for Evaluating Reproducibility of LLMs in Data Science</i>	
Qiuhai Zeng, Claire Jin, Xinyue Wang, Yuhan Zheng and Qunhua Li	10170
<i>MisinfoBench: A Multi-Dimensional Benchmark for Evaluating LLMs' Resilience to Misinformation</i>	
Ye Yang, Donghe Li, Zuchen Li, Fengyuan Li, Jingyi Liu, Li Sun and Qingyu Yang	10202
<i>Fuzzy Reasoning Chain (FRC): An Innovative Reasoning Framework from Fuzziness to Clarity</i>	
Ping Chen, Xiang Liu, Zhaoxiang Liu, Zezhou Chen, Xingpeng Zhang, Huan Hu, Zipeng Wang, Kai Wang, Shuming Shi and Shiguo Lian	10230
<i>HighMATH: Evaluating Math Reasoning of Large Language Models in Breadth and Depth</i>	
Yan Liu, Minghui Zhang, Bojian Xiong, Yifan Xiao, Yinong Sun, Yating Mei, Longyu Zeng, Jingchao Yang, Yang Wang and Deyi Xiong	10241
<i>CATCH: A Novel Data Synthesis Framework for High Therapy Fidelity and Memory-Driven Planning Chain of Thought in AI Counseling</i>	
Mingyu Chen, Jingkai Lin, Zhaojie Chu, Xiaofen Xing, Yirong Chen and Xiangmin Xu	10254
<i>MediVLM: A Vision Language Model for Radiology Report Generation from Medical Images</i>	
Debanjan Goswami, Ronast Subedi and Shayok Chakraborty	10287
<i>AdDriftBench: A Benchmark for Detecting Data Drift and Label Drift in Short Video Advertising</i>	
Yinghao Song, Xiangji Zeng, Shuai Cui, Lu Sun, Zhaowei Liu, Yuan Yuan, Yulu Wang, Hai Zhou and Zhaohan Gong	10305
<i>NIM: Neuro-symbolic Ideographic Metalanguage for Inclusive Communication</i>	
Prawaal Sharma, Poonam Goyal, Navneet Goyal and Vidisha Sharma	10322
<i>ViFT: Towards Visual Instruction-Free Fine-tuning for Large Vision-Language Models</i>	
Zikang Liu, Kun Zhou, Xin Zhao, Dawei Gao, Yaliang Li and Ji-Rong Wen	10341

<i>Do Code Semantics Help? A Comprehensive Study on Execution Trace-Based Information for Code Large Language Models</i>	
Jian Jornbowl Wang, Xiaofei Xie, Qiang Hu, Shangqing Liu and Yi Li	10367
<i>LongWeave: A Long-Form Generation Benchmark Bridging Real-World Relevance and Verifiability</i>	
Zikai Xiao, Fei Huang, Jianhong Tu, Jianhui Wei, Wen MA, Yuxuan Zhou, Jian Wu, Bowen Yu, Zuozhu Liu and Junyang Lin	10386
<i>XL-Suite: Cross-Lingual Synthetic Training and Evaluation Data for Open-Ended Generation</i>	
Vivek Iyer, Pinzhen Chen, Ricardo Rei and Alexandra Birch	10418
<i>Accelerating LLM Reasoning via Early Rejection with Partial Reward Modeling</i>	
Seyyed Saeid Cheshmi, Azal Ahmad Khan, Xinran Wang, Zirui Liu and Ali Anwar	10433
<i>CultureSynth: A Hierarchical Taxonomy-Guided and Retrieval-Augmented Framework for Cultural Question-Answer Synthesis</i>	
Xinyu Zhang, Pei Zhang, Shuang Luo, Jialong Tang, Yu Wan, Baosong Yang and Fei Huang	10448
<i>DesignCLIP: Multimodal Learning with CLIP for Design Patent Understanding</i>	
Zhu Wang, Homaira Huda Shomee, Sathya N. Ravi and Sourav Medya	10468
<i>R3-RAG: Learning Step-by-Step Reasoning and Retrieval for LLMs via Reinforcement Learning</i>	
Yuan Li, Qi Luo, Xiaonan Li, Bufan Li, Qinyuan Cheng, Bo Wang, Yining Zheng, Yuxin Wang, Zhangyue Yin and Xipeng Qiu	10491
<i>'Hello, World!': Making GNNs Talk with LLMs</i>	
Sunwoo Kim, Soo Yong Lee, Jaemin Yoo and Kijung Shin	10508
<i>Both Text and Images Leaked! A Systematic Analysis of Data Contamination in Multimodal LLM</i>	
Dingjie Song, Sicheng Lai, Mingxuan Wang, Shunian Chen, Lichao Sun and Benyou Wang	10527
<i>NLKI: A Lightweight Natural Language Knowledge Integration Framework for Improving Small VLMs in Commonsense VQA Tasks</i>	
Aritra Dutta, Swapnanil Mukherjee, Deepanway Ghosal and Somak Aditya	10543
<i>Text or Pixels? Evaluating Efficiency and Understanding of LLMs with Visual Text Inputs</i>	
Yanhong Li, Zixuan Lan and Jiawei Zhou	10564
<i>Assessing Socio-Cultural Alignment and Technical Safety of Sovereign LLMs</i>	
Kyubung Chae, Gihoon Kim, Gyuseong Lee, Taesup Kim, Jaejin Lee and Heejin Kim ...	10579
<i>Sample Efficient Alignment Learning With Episodic Control</i>	
Van Dai Do, Quan Hung Tran, Ahmed Kirmani, Lu Zhang and Hung Le	10601
<i>Evaluating Automatic Speech Recognition Systems for Korean Meteorological Experts</i>	
ChaeHun Park, Hojun Cho and Jaegul Choo	10619
<i>3D-Aware Vision-Language Models Fine-Tuning with Geometric Distillation</i>	
Seonho Lee, Jiho Choi, Inha Kang, Jiwook Kim, Junsung Park and Hyunjung Shim	10628
<i>CAPE: Context-Aware Personality Evaluation Framework for Large Language Models</i>	
Jivnesh Sandhan, Fei Cheng, Tushar Sandhan and Yugo Murawaki	10648
<i>AgentThink: A Unified Framework for Tool-Augmented Chain-of-Thought Reasoning in Vision-Language Models for Autonomous Driving</i>	
Kangan Qian, Sicong Jiang, Yang Zhong, Ziang Luo, Zilin Huang, Tianze Zhu, Kun Jiang, Mengmeng Yang, Zheng Fu, Jinyu Miao, Yining Shi, He Zhe Lim, Li Liu, Tianbao Zhou, Hongyi Wang, Huang Yu, Yifei HU, Guang Li, Guang Chen, Hao Ye, Lijun Sun and Diange Yang	10663

<i>Select to Know: An Internal-External Knowledge Self-Selection Framework for Domain-Specific Question Answering</i>	
Bolei He, Xinran He, Run Shao, Shanfu Shu, Xianwei Xue, MingQuan Cheng, Haifeng Li and Zhen-Hua Ling	10683
<i>GenPTQ: Green Post-Training Quantization for Large-Scale ASR Models with Mixed-Precision Bit Allocation</i>	
Beom Jin Kang and Hyun Kim	10704
<i>Where Does This Strange Smell Come from?": Enabling Conversational Interfaces for Artificial Olfaction</i>	
Xueyi Zhou, Qi Lu and Dong-Kyu Chae	10719
<i>LightRAG: Simple and Fast Retrieval-Augmented Generation</i>	
Zirui Guo, Lianghao Xia, Yanhua Yu, Tu Ao and Chao Huang	10746
<i>Beyond Distribution: Investigating Language Models' Understanding of Sino-Korean Morphemes</i>	
Taehee Jeon	10762
<i>Sarcasm-RI: Enhancing Sarcasm Detection through Focused Reasoning</i>	
Qi Yang, Jingjie Zeng, Liang Yang, Kai Ma and Hongfei Lin	10773
<i>ISACL: Internal State Analyzer for Copyrighted Training Data Leakage</i>	
Guangwei Zhang, Qisheng Su, Jiateng Liu, Cheng Qian, Yanzhou Pan, Yanjie Fu and Denghui Zhang	10786
<i>Steering LVLMS via Sparse Autoencoder for Hallucination Mitigation</i>	
Zhenglin Hua, Jinghan He, Zijun Yao, Tianxu Han, Haiyun Guo, Yuheng Jia and Junfeng Fang	10808
<i>On the Perception Bottleneck of VLMs for Chart Understanding</i>	
Junteng Liu, Weihao Zeng, Xiwen Zhang, Yijun Wang, Zifei Shan and Junxian He	10829
<i>Self-Guided Function Calling in Large Language Models via Stepwise Experience Recall</i>	
Sijia Cui, Aiyao He, Shuai Xu, Hongming Zhang, Yanna Wang, Qingyang Zhang, Yajing Wang and bo xu	10842
<i>Multilingual Generative Retrieval via Cross-lingual Semantic Compression</i>	
Yuxin Huang, Simeng Wu, Ran Song, Yan Xiang, Yantuan Xian, Shengxiang Gao and Zhengtao Yu	10855
<i>Towards Multi-Document Question Answering in Scientific Literature: Pipeline, Dataset, and Evaluation</i>	
Hui Huang, Julien Velcin and Yacine Kessaci	10867
<i>Multilingual Knowledge Graph Completion via Efficient Multilingual Knowledge Sharing</i>	
Cunli Mao, Xiaofei Gao, Ran Song, Shizhu He, Shengxiang Gao, Kang Liu and Zhengtao Yu	10882
<i>Mitigating Attention Localization in Small Scale: Self-Attention Refinement via One-step Belief Propagation</i>	
Nakyung Lee, Yeongoon Kim, Minhae Oh, Suhwan Kim, Jin Woo Koo, Hyewon Jo and Jungwoo Lee	10897
<i>Imagination and Contemplation: A Balanced Framework for Semantic-Augmented Multimodal Machine Translation</i>	
Zhuang Yu, Shiliang Sun, Jing Zhao, Tengfei Song and Hao Yang	10913

<i>NeLLCom-Lex: A Neural-agent Framework to Study the Interplay between Lexical Systems and Language Use</i>	
Yuqing Zhang, Ecesu Ürker, Tessa Verhoef, Gemma Boleda and Arianna Bisazza	10929
<i>RLMEval: Evaluating Research-Level Neural Theorem Proving</i>	
Auguste Poiroux, Antoine Bosselut and Viktor Kunčák	10946
<i>KaeDe: Progressive Generation of Logical Forms via Knowledge-Aware Question Decomposition for Improved KBQA</i>	
Ranran Bu, Jian Cao, Jianqi Gao, Shiyu Qian and Hongming Cai	10958
<i>Where Fact Ends and Fairness Begins: Redefining AI Bias Evaluation through Cognitive Biases</i>	
Jen-tse Huang, Yuhang Yan, Linqi Liu, Yixin Wan, Wenxuan Wang, Kai-Wei Chang and Michael R. Lyu	10974
<i>Equal Truth: Rumor Detection with Invariant Group Fairness</i>	
Junyi Chen, Mengjia Wu, Qian Liu, Jing Sun, Ying Ding and Yi Zhang	10994
<i>STEAM: A Semantic-Level Knowledge Editing Framework for Large Language Models</i>	
Geunyeong Jeong, Juoh Sun, Seonghee Lee and Harksoo Kim	11008
<i>SoT: Structured-of-Thought Prompting Guides Multilingual Reasoning in Large Language Models</i>	
Rui Qi, Zhibo Man, Yufeng Chen, Fengran Mo, Jinan Xu and Kaiyu Huang	11024
<i>How Reliable is Multilingual LLM-as-a-Judge?</i>	
Xiyan Fu and Wei Liu	11040
<i>Cognitive-Level Adaptive Generation via Capability-Aware Retrieval and Style Adaptation</i>	
Qingsong Wang, Tao Wu, Wang Lin, Yueying Feng, Gongsheng Yuan, Chang Yao and Jingyuan Chen	11054
<i>Data Doping or True Intelligence? Evaluating the Transferability of Injected Knowledge in LLMs</i>	
Essa Jan, Moiz Ali, Muhammad Saram Hassan, Muhammad Fareed Zaffar and Yasir Zaki	11070
<i>INDOORWORLD : Integrating Physical Task Solving and Social Simulation in A Heterogeneous Multi-Agent Environment</i>	
Dekun Wu, Frederik Brudy, Bang Liu and Yi Wang	11078
<i>ARXSA: A General Negative Feedback Control Theory in Vision-Language Models</i>	
Zeyu Zhang, Tianqi Chen and Yuki Todo	11100
<i>Breaking the Attention Trap in Code LLMs: A Rejection Sampling Approach to Enhance Code Execution Prediction</i>	
Xingcheng Ruan, Haoxiang Geng, Yunhui Xia and Bingran Zhao	11111
<i>HiMATE: A Hierarchical Multi-Agent Framework for Machine Translation Evaluation</i>	
Shijie Zhang, Renhao Li, Songsheng Wang, Philipp Koehn, Min Yang and Derek F. Wong	11121
<i>ReliableEval: A Recipe for Stochastic LLM Evaluation via Method of Moments</i>	
Gili Lior, Eliya Habba, Shahar Levy, Avi Caciularu and Gabriel Stanovsky	11146
<i>From Characters to Tokens: Dynamic Grouping with Hierarchical BPE</i>	
Rares Dolga, Lucas Maystre, Tudor Berariu and David Barber	11154
<i>Auto-SLURP: A Benchmark Dataset for Evaluating Multi-Agent Frameworks in Smart Personal Assistant</i>	
Lei Shen and Xiaoyu Shen	11163

<i>NER Retriever: Zero-Shot Named Entity Retrieval with Type-Aware Embeddings</i> Or Shachar, Uri Katz, Yoav Goldberg and Oren Glickman	11175
<i>MMATH: A Multilingual Benchmark for Mathematical Reasoning</i> Wenyang Luo, Xin Zhao, Jing Sha, Shijin Wang and Ji-Rong Wen	11187
<i>MultiClaimNet: A Massively Multilingual Dataset of Fact-Checked Claim Clusters</i> Rubaa Panchendrarajan, Rubén Míguez Pérez and Arkaitz Zubiaga	11203
<i>DS-MHP: Improving Chain-of-Thought through Dynamic Subgraph-Guided Multi-Hop Path</i> Yongqiang Liu, Qiyao Peng, Binrong Liu, Hongtao Liu, XueWei Li and Wenjun Wang ...	11216
<i>LongTail-Swap: benchmarking language models' abilities on rare words</i> Robin Algayres, Charles-Éric Saint-James, Mahi Luthra, Jiayi Shen, Youssef Bencheekroun, Dongyan Lin, Rashel Moritz, Juan Pino and Emmanuel Dupoux	11231
<i>TF-Mamba: Text-enhanced Fusion Mamba with Missing Modalities for Robust Multimodal Sentiment Analysis</i> Xiang Li, Xianfu Cheng, Dezhuang Miao, Xiaoming Zhang and Zhoujun Li	11252
<i>Are Economists Always More Introverted? Analyzing Consistency in Persona-Assigned LLMs</i> Manon Reusens, Bart Baesens and David Jurgens	11268
<i>Can you SPLICE it together? A Human Curated Benchmark for Probing Visual Reasoning in VLMs</i> Mohamad Ballout, Okajevo Wilfred, Seyedalireza Yaghoubi, Nohayr Muhammad Abdelmoneim, Julius Mayer and Elia Bruni	11288
<i>On the Effectiveness of Prompt-Moderated LLMs for Math Tutoring at the Tertiary Level</i> Sebastian Steindl, Fabian Brunner, Nada Sissouno, Dominik Schwagerl, Florian Schöler-Niewiera and Ulrich Schäfer	11310
<i>SkewRoute: Training-Free LLM Routing for Knowledge Graph Retrieval-Augmented Generation via Score Skewness of Retrieved Context</i> Hairu Wang, Yuan Feng, Yukun Cao, Xike Xie and S Kevin Zhou	11324
<i>Acquiescence Bias in Large Language Models</i> Daniel Braun	11341
<i>Time to Talk: LLM Agents for Asynchronous Group Communication in Mafia Games</i> Niv Eckhaus, Uri Berger and Gabriel Stanovsky	11356
<i>How Sampling Affects the Detectability of Machine-written texts: A Comprehensive Study</i> Matthieu Dubois, François Yvon and Pablo Piantanida	11369
<i>An Improved, Strong Baseline for Pre-Trained Large Language Models as Task-Oriented Dialogue Systems</i> Sebastian Steindl, André Kestler, Ulrich Schäfer and Bernd Ludwig	11388
<i>MATCH: Task-Driven Code Evaluation through Contrastive Learning</i> Marah Ghoummaid, Vladimir Tchuiev, Ofek Glick, Michal Moshkovitz and Dotan Di Castro	11399
<i>Evaluating Large Language Models for Cross-Lingual Retrieval</i> Longfei Zuo, Pingjun Hong, Oliver Kraus, Barbara Plank and Robert Litschko	11415
<i>SGCD: Subtask-Guided Causal-Debiasing Framework for Robust Cross-Utterance Sentiment Quadruple Extraction in Dialogues</i> Xiang Li, Keyu Yao and Gang Shen	11430

<i>FaMTEB: Massive Text Embedding Benchmark in Persian Language</i>	
Erfan Zinvandi, Morteza Alikhani, Mehran Sarmadi, Zahra Pourbahman, Sepehr Arvin, Reza Kazemi and Arash Amini	11441
<i>Leveraging High-Resource English Corpora for Cross-lingual Domain Adaptation in Low-Resource Japanese Medicine via Continued Pre-training</i>	
Kazuma Kobayashi, Zhen Wan, Fei Cheng, Yuma Tsuta, Xin Zhao, Junfeng Jiang, Jiahao Huang, Zhiyi Huang, Yusuke Oda, Rio Yokota, Yuki Arase, Daisuke Kawahara, Akiko Aizawa and Sadao Kurohashi	11469
<i>Structure Trumps Size: Rethinking Data Quality for LLM Reasoning</i>	
Hu Xu, Zeyan Li, Rui Wang and Jianfeng Xu	11489
<i>A Zero-Shot Neuro-Symbolic Approach for Complex Knowledge Graph Question Answering</i>	
Perna Agarwal and Srikanta Bedathur	11514
<i>Making Every Step Effective: Jailbreaking Large Vision-Language Models Through Hierarchical KV Equalization</i>	
Shuyang Hao, Yiwei Wang, Bryan Hooi, Jun Liu, Muhao Chen, Zi Huang and Yujun Cai	11528
<i>MT-Mol: Multi Agent System with Tool-based Reasoning for Molecular Optimization</i>	
Hyomin Kim, Yunhui Jang and Sungsoo Ahn	11544
<i>A Survey on LLM-powered Agents for Recommender Systems</i>	
Qiyao Peng, Hongtao Liu, Hua Huang, Jian Yang, Qing Yang and Minglai Shao	11574
<i>Efficiently Selecting Response Generation Strategies for Synthetic Data Construction by Self-Aligned Perplexity</i>	
Xuan Ren, Qi Chen and Lingqiao Liu	11584
<i>Benchmarking for Domain-Specific LLMs: A Case Study on Academia and Beyond</i>	
Rubing Chen, Jiaxin Wu, Jian Wang, Xulu Zhang, Wenqi Fan, Chenghua Lin, Xiaoyong Wei and Li Qing	11606
<i>FrameEOL: Semantic Frame Induction using Causal Language Models</i>	
Chihiro Yano, Kosuke Yamada, Hayato Tsukagoshi, Ryohei Sasano and Koichi Takeda	11620
<i>CaTER: A Framework for Context-aware Topology Entity Retrieval Contrastive Learning in End-to-End Task-Oriented Dialogue Systems</i>	
Di Wu Hebeu and Zhizhi Yu	11633
<i>Attribution and Application of Multiple Neurons in Multimodal Large Language Models</i>	
Feiyu Wang, Ziran Zhao, Dong Yu and Pengyuan Liu	11649
<i>When Models Lie, We Learn: Multilingual Span-Level Hallucination Detection with PsiloQA</i>	
Elisei Rykov, Kseniia Petrushina, Maksim Savkin, Valerii Olisov, Artem Vazhentsev, Kseniia Titova, Alexander Panchenko, Vasily Konovalov and Julia Belikova	11663
<i>Unraveling Misinformation Propagation in LLM Reasoning</i>	
Yiyang Feng, Yichen Wang, Shaobo Cui, Boi Faltings, Mina Lee and Jiawei Zhou	11683
<i>RAISE: Reinforced Adaptive Instruction Selection For Large Language Models</i>	
Qingsong Lv, Yangning Li, Zihua Lan, Zishan Xu, Jiwei Tang, Tingwei Lu, Yinghui Li, Wenhao Jiang, Hong-Gee Kim, Hai-Tao Zheng and Philip S. Yu	11708
<i>Teaching According to Talents! Instruction Tuning LLMs with Competence-Aware Curriculum Learning</i>	
Yangning Li, Tingwei Lu, Yinghui Li, Yankai Chen, Wei-Chieh Huang, Wenhao Jiang, Hui Wang, Hai-Tao Zheng and Philip S. Yu	11724

<i>Let Them Down Easy! Contextual Effects of LLM Guardrails on User Perceptions and Preferences</i> Mingqian Zheng, Wenjia hu, Patrick Zhao, Motahhare Eslami, Jena D. Hwang, Faeze Brahman, Carolyn Rose and Maarten Sap	11742
<i>From Hypothesis to Publication: A Comprehensive Survey of AI-Driven Research Support Systems</i> Zekun Zhou, Xiaocheng Feng, Lei Huang, Xiachong Feng, Ziyun Song, Ruihan Chen, Liang Zhao, Weitao Ma, Yuxuan Gu, Baoxin Wang, Dayong Wu, Guoping Hu, Ting Liu and Bing Qin	11773
<i>Enhancing Model Privacy in Federated Learning with Random Masking and Quantization</i> Zhibo Xu, Zhu JianHao, Jingwen Xu, Changze Lv, Zhenghua Wang, Zisu Huang, Xiaohua Wang, Muling Wu, Qi Qian, Xiaoqing Zheng and Xuanjing Huang	11804
<i>SuPreME: A Supervised Pre-training Framework for Multimodal ECG Representation Learning</i> Mingsheng Cai, Jiuming Jiang, Wenhao Huang, Che Liu and Rossella Arcucci	11817
<i>Ferret: Faster and Effective Automated Red Teaming with Reward-Based Scoring Technique</i> Tej Deep Pala, Vernon Toh, Rishabh Bhardwaj and Soujanya Poria	11845
<i>Do What? Teaching Vision-Language-Action Models to Reject the Impossible</i> Wen-Han Hsieh, Elvis Hsieh, Dantong Niu, Trevor Darrell, Roei Herzig and David M. Chan	11861
<i>AgentInit: Initializing LLM-based Multi-Agent Systems via Diversity and Expertise Orchestration for Effective and Efficient Collaboration</i> Chunhao Tian, Yutong Wang, Xuebo Liu, Zhexuan Wang, Liang Ding, Miao Zhang and Min Zhang	11870
<i>Time to Revisit Exact Match</i> Auss Abbood, Zaiqiao Meng and Nigel Collier	11903
<i>LongTableBench: Benchmarking Long-Context Table Reasoning across Real-World Formats and Do- mains</i> Liyao Li, Jiaming Tian, Hao Chen, Wentao Ye, Chao Ye, Haobo Wang, Ningtao Wang, Xing Fu, Gang Chen and Junbo Zhao	11927
<i>Exploring and Evaluating Multimodal Knowledge Reasoning Consistency of Multimodal Large Lan- guage Models</i> Boyue Jia, Junzhe Zhang, Huixuan Zhang and Xiaojun Wan	11966
<i>MPTA: MultiTask Personalization Assessment</i> Matthieu Tehenan, Eric Chamoun and Andreas Vlachos	11982
<i>Semantic Geometry of Sentence Embeddings</i> Matthieu Tehenan	11993
<i>ReAlign: Structured Revision for Small Language Model Alignment</i> Ruijun Chen, Jiajian Guo, Hongzhan Chen, Fanqi Wan, Qifan Wang and Xiaojun Quan ...	12005
<i>Curr-ReFT: Overcoming Training Bottlenecks in Small-scale Vision-Language Models via Curriculum Reinforcement Finetuning</i> Huilin Deng, Ding Zou, Xinghao Zhao, Rui Ma, Yanming Guo, Yang Cao and Yu Kang ..	12021
<i>Layer-Aware Task Arithmetic: Disentangling Task-Specific and Instruction-Following Knowledge</i> Yan-Lun Chen, Yi-Ru Wei, Chia-Yi Hsu, Chia-Mu Yu, Chun-Ying Huang, Ying-Dar Lin, Yu-Sung Wu and Wei-Bin Lee	12033
<i>Revisiting Pruning vs Quantization for Small Language Models</i> Zihan Zhou, Simon Kurz and Zhixue Zhao	12055

<i>CLaw: Benchmarking Chinese Legal Knowledge in Large Language Models - A Fine-grained Corpus and Reasoning Analysis</i>	
Xinzhe Xu, Liang Zhao, Hongshen Xu and Chenchenc	12071
<i>polyBART: A Chemical Linguist for Polymer Property Prediction and Generative Design</i>	
Anagha Savit, Harikrishna Sahu, Shivank S. Shukla, Wei Xiong and Rampi Ramprasad	12104
<i>A Survey of RAG-Reasoning Systems in Large Language Models</i>	
Yangning Li, Weizhi Zhang, Yuyao Yang, Wei-Chieh Huang, Yaozu Wu, Junyu Luo, Yuanchen Bei, Henry Peng Zou, Xiao Luo, Yusheng Zhao, Chunkit Chan, Yankai Chen, Zhongfen Deng, Yinghui Li, Hai-Tao Zheng, Dongyuan Li, Renhe Jiang, Ming Zhang, Yangqiu Song and Philip S. Yu	12120
<i>REGen: A Reliable Evaluation Framework for Generative Event Argument Extraction</i>	
Omar Sharif, Joseph Gatto, Madhusudan Basak and Sarah Masud Preum	12146
<i>Mitigating Interviewer Bias in Multimodal Depression Detection: An Approach with Adversarial Learning and Contextual Positional Encoding</i>	
Enshi Zhang and Christian Poellabauer	12169
<i>AMIA: Automatic Masking and Joint Intention Analysis Makes LVLMS Robust Jailbreak Defenders</i>	
Yuqi Zhang, Yuchun Miao, Zuchao Li and Liang Ding	12189
<i>Disentangling Language Understanding and Reasoning Structures in Cross-lingual Chain-of-Thought Prompting</i>	
Khanh-Tung Tran, Nguyet-Hang Vu, Barry O’Sullivan and Hoang D. Nguyen	12200
<i>MoRoVoc: A Large Dataset for Geographical Variation Identification of the Spoken Romanian Language</i>	
Andrei-Marius Avram, Bănescu Ema-Ioana, Anda-Teodora Robea, Dumitru-Clementin Cercel and Mihaela-Claudia Cercel	12207
<i>Language-Informed Synthesis of Rational Agent Models for Grounded Theory-of-Mind Reasoning On-the-fly</i>	
Lance Ying, Ryan Truong, Katherine M. Collins, Cedegao E. Zhang, Megan Wei, Tyler Brooke Wilson, Tan Zhi-Xuan, Lionel Wong and Joshua B. Tenenbaum	12217
<i>MOLE: Metadata Extraction and Validation in Scientific Papers Using LLMs</i>	
Zaid Alyafeai, Maged S. Al-shaibani and Bernard Ghanem	12236
<i>MASSV: Multimodal Adaptation and Self-Data Distillation for Speculative Decoding of Vision-Language Models</i>	
Mugilan Ganesan, Shane Segal, Ankur Aggarwal, Nish Sinnadurai, Sean Lie and Vithursan Thangarasa	12265
<i>FESTA: Functionally Equivalent Sampling for Trust Assessment of Multimodal LLMs</i>	
Debarpan Bhattacharya, Apoorva Kulkarni and Sriram Ganapathy	12277
<i>ClaimGen-CN: A Large-scale Chinese Dataset for Legal Claim Generation</i>	
Siying Zhou, Yiquan Wu, Hui Chen, Xueyu Hu, Kun Kuang, Adam Jatowt, Chunyan Zheng and Fei Wu	12296
<i>Summarize-Exemplify-Reflect: Data-driven Insight Distillation Empowers LLMs for Few-shot Tabular Classification</i>	
Yifei Yuan, Jiatong Li, Weijia Zhang, Mohammad Aliannejadi, Evangelos Kanoulas and Renjun Hu	12324

<i>Rethinking LLM Uncertainty: A Multi-Agent Approach to Estimating Black-Box Model Uncertainty</i> Yu Feng, Phu Mon Htut, Zheng Qi, Wei Xiao, Manuel Mager, Nikolaos Pappas, Kishaloy Halder, Yang Li, Yassine Benajiba and Dan Roth	12349
<i>Stress-Testing the Reasoning Competence of Language Models With Formal Proofs</i> Konstantine Arkoudas and Serafim Batzoglou	12376
<i>Topic-Guided Reinforcement Learning with LLMs for Enhancing Multi-Document Summarization</i> Chuyuan Li, Austin Xu, Shafiq Joty and Giuseppe Carenini	12395
<i>FACTCHECKMATE: Preemptively Detecting and Mitigating Hallucinations in LMs</i> Deema Alnuhait, Neeraja Kirtane, Muhammad Khalifa and Hao Peng	12413
<i>Dialectal Toxicity Detection: Evaluating LLM-as-a-Judge Consistency Across Language Varieties</i> Fahim Faisal, Md Mushfiqur Rahman and Antonios Anastasopoulos	12429
<i>Mitigate One, Skew Another? Tackling Intersectional Biases in Text-to-Image Models</i> Pushkar Shukla, Aditya Chinchure, Emily Diana, Alexander Tolbert, Kartik Hosanagar, Vineeth N. Balasubramanian, Leonid Sigal and Matthew A. Turk	12453
<i>Language-Specific Layer Matters: Efficient Multilingual Enhancement for Large Vision-Language Mo- dels</i> Yuchun Fan, Yilin Wang, Yongyu Mu, Lei Huang, Bei Li, Xiaocheng Feng, Tong Xiao and JingBo Zhu	12473
<i>InfAL: Inference Time Adversarial Learning for Improving Research Ideation</i> Sikun Guo, Amir Hassan Shariatmadari, Peng Wang, Albert Huang and Aidong Zhang . . .	12501
<i>Speculative Decoding for Multi-Sample Inference</i> Yiwei Li, Jiayi Shi, Shaoxiong Feng, Peiwen Yuan, Xinglin Wang, Yueqi Zhang, Ji Zhang, Chuyi Tan, Boyuan Pan, Yao Hu and Kan Li	12523
<i>LSRL: Process-Supervised GRPO on Latent Recurrent States Improves Mathematical Reasoning</i> Hangliang Ren	12534
<i>Multi-token Mask-filling and Implicit Discourse Relations</i> Meinan Liu, Yunfang Dong, Xixian Liao and Bonnie Webber	12546
<i>Schema Generation for Large Knowledge Graphs Using Large Language Models</i> Bohui Zhang, Yuan He, Lydia Pintscher, Albert Meroño-Peñuela and Elena Simperl	12561
<i>MCTS-RAG: Enhancing Retrieval-Augmented Generation with Monte Carlo Tree Search</i> Yunhai Hu, Yilun Zhao, Chen Zhao and Arman Cohan	12581
<i>What if Othello-Playing Language Models Could See?</i> Xinyi Chen, Yifei Yuan, Jiaang Li, Serge Belongie, Maarten de Rijke and Anders Søgaard	12598
<i>LLM-Based Web Data Collection for Research Dataset Creation</i> Thomas Berkane, Marie-Laure Charpignon and Maimuna S. Majumder	12610
<i>PsyScam: A Benchmark for Psychological Techniques in Real-World Scams</i> Shang Ma, Tianyi Ma, Jiahao Liu, Wei Song, Zhenkai Liang, Xusheng Xiao and Yanfang Ye	12623
<i>LoRaDA: Low-Rank Direct Attention Adaptation for Efficient LLM Fine-tuning</i> Zhangming Li, Qinghao Hu, Yiqun Chen, Peisong Wang, Yifan Zhang and Jian Cheng . . .	12638

<i>Inductive Reasoning on Few-Shot Knowledge Graphs with Task-Aware Language Models</i> Cheng Yan, Feng Zhao, Ruilin Zhao and Hong Zhang	12656
<i>ForestCast: Open-Ended Event Forecasting with Semantic News Forest</i> Zi Yu, Shaoxiang Wang, Guozheng Li, Yu Zhang and Chi Harold Liu	12667
<i>Agentic Medical Knowledge Graphs Enhance Medical Question Answering: Bridging the Gap Between LLMs and Evolving Medical Knowledge</i> Mohammad Reza Rezaei, Reza Saadati Fard, Jayson Lee Parker, Rahul G Krishnan and Milad Lankarany	12682
<i>Text Anomaly Detection with Simplified Isolation Kernel</i> Yang Cao, Sikun Yang, Yujiu Yang, Lianyong Qi and Ming Liu	12702
<i>Idola Tribus of AI: Large Language Models tend to perceive order where none exists</i> Shin-nosuke Ishikawa, Masato Todo, Taiki Ogihara and Hirotsugu Ohba	12714
<i>Thunder-DeID: Accurate and Efficient De-identification Framework for Korean Court Judgments</i> Sungeun Hahm, Heejin Kim, Gyuseong Lee, Hyunji M. Park and Jaejin Lee	12728
<i>Multi-Agent Autonomous Driving Systems with Large Language Models: A Survey of Recent Advances, Resources, and Future Directions</i> Yaozu Wu, Dongyuan Li, Yankai Chen, Renhe Jiang, Henry Peng Zou, Wei-Chieh Huang, Yan-ning Li, Liancheng Fang, Zhen Wang and Philip S. Yu	12756
<i>Comprehensive Evaluation on Lexical Normalization: Boundary-Aware Approaches for Unsegmented Languages</i> Shohei Higashiyama and Masao Utiyama	12774
<i>Explainable Text Classification with LLMs: Enhancing Performance through Dialectical Prompting and Explanation-Guided Training</i> Huaming Du, Lei Yuan, Cancan Feng, Guisong Liu, Gang Kou and Carl Yang	12800
<i>MultiPL-MoE: Multi-Programming-Lingual Extension of Large Language Models through Hybrid Mixture-of-Experts</i> Qing Wang, Xue Han, Jiahui Wang, Lehao Xing, Qian Hu, Lianlian Zhang, Chao Deng and Junlan Feng	12817
<i>AutoSpec: An Agentic Framework for Automatically Drafting Patent Specification</i> Ryan Shea and Zhou Yu	12829
<i>LimaCost: Data Valuation for Instruction Tuning of Large Language Models</i> Hyeonseok Moon, Jaehyung Seo, Seonmin Koo, Jinsung Kim, Young-kyoung Ham, Jiwon Moon and Heuseok Lim	12841
<i>Two Challenges, One Solution: Robust Multimodal Learning through Dynamic Modality Recognition and Enhancement</i> Lanxin Bi, Yunqi Zhang, Luyi Wang, Yake Niu and Hui Zhao	12855
<i>SwiftPrune: Hessian-Free Weight Pruning for Large Language Models</i> Yuhan Kang, Yang Shi, Mei Wen, Jun He, Jianchao Yang, Zeyu Xue, Jing Feng and Xinwang Liu	12868
<i>Training LLMs for Optimization Modeling via Iterative Data Synthesis and Structured Validation</i> Yang Wu, Yifan Zhang, Yurong Wu, Yuran Wang, Junkai Zhang and Jian Cheng	12880

<i>Exploiting Prompt-induced Confidence for Black-Box Attacks on LLMs</i>	
Meina Chen, Yihong Tang and Kehai Chen	12897
<i>DPF-CM: A Data Processing Framework with Privacy-Preserving Vector Databases for Chinese Medical LLMs Training and Deployment</i>	
Wei Huang, Anda Cheng, Zhao Zhang and Yinggui Wang	12904
<i>Graph-Reward-SQL: Execution-Free Reinforcement Learning for Text-to-SQL via Graph Matching and Stepwise Reward</i>	
Han Weng, Puzhen Wu, Cui Longjie, Yi Zhan, Boyi Liu, Yuanfeng Song, Dun Zeng, Yingxiang Yang, Qianru Zhang, Dong Huang, Xiaoming Yin, Yang Sun and Xing Chen	12917
<i>StatsChartMWP: A Dataset for Evaluating Multimodal Mathematical Reasoning Abilities on Math Word Problems with Statistical Charts</i>	
Dan Zhu, Tianqiao Liu and Zitao Liu	12944
<i>Logic-Thinker: Teaching Large Language Models to Think more Logically.</i>	
Chengyao Wen, Qiang Cheng, Shaofei Wang, Zhizhen Liu, Deng Zhao and Lei Liang	12955
<i>ACEBench: A Comprehensive Evaluation of LLM Tool Usage</i>	
Chen Chen, Xinlong Hao, Weiwen Liu, Xu Huang, Xingshan Zeng, Shuai Yu, Dexun Li, Yuefeng Huang, Xiangcheng Liu, Wang Xinzhi and Wu Liu	12970
<i>RevPRAG: Revealing Poisoning Attacks in Retrieval-Augmented Generation through LLM Activation Analysis</i>	
Xue Tan, Hao Luan, Mingyu Luo, Xiaoyan Sun, Ping Chen and Jun Dai	12999
<i>DaMoC: Efficiently Selecting the Optimal Large Language Model for Fine-tuning Domain Tasks Based on Data and Model Compression</i>	
Wei Huang, Huang Wei and Yinggui Wang	13012
<i>CoAT: Chain-of-Associated-Thoughts Framework for Enhancing Large Language Models Reasoning</i>	
Jianfeng Pan, Senyou Deng and Shaomang Huang	13028
<i>ChartM³: A Multi-Stage Code-Driven Pipeline for Constructing Multi-Dimensional and Multi-Step Visual Reasoning Data in Chart Comprehension</i>	
Duo Xu, Hao Cheng, Xin Lin, Zhen Xie and Hao Henry Wang	13046
<i>Can LLMs Truly Plan? A Comprehensive Evaluation of Planning Capabilities</i>	
Gayeon Jung, HyeonSeok Lim, Minjun Kim, Joon-ho Lim, KyungTae Lim and Hansaem Kim	13069
<i>MARIO-0.5B: A Multi-Agent Lightweight Model for Real-Time Open Information Extraction in Low-Resource Settings</i>	
Donghai Zhang, SHuangtao Yang, Dong Xiaozheng, Wei Song and Bo Fu	13085
<i>BiMax: Bidirectional MaxSim Score for Document-Level Alignment</i>	
Xiaotian Wang, Takehito Utsuro and Masaaki Nagata	13095
<i>DocMMIR: A Framework for Document Multi-modal Information Retrieval</i>	
Zirui Li, Siwei Wu, Yizhi LI, Xingyu Wang, Yi Zhou and Chenghua Lin	13117
<i>MoVoC: Morphology-Aware Subword Construction for Ge'ez Script Languages</i>	
Hailay Kidu Teklehaymanot, Dren Fazlija and Wolfgang Nejdl	13131
<i>MMA: Cross-Domain Knowledge Integration via Mixture of Multi-Domain Agents</i>	
Kehang Jia, Juntao Li, Xiaobo Liang, Yisheng Xiao, Yixuan Yang and Min Zhang	13145

<i>HAWK: Highlighting Entity-aware Knowledge for Alleviating Information Sparsity in Long Contexts</i> Seonmin Koo, Jinsung Kim, Chanjun Park and Heuseok Lim	13161
<i>Sensitivity-LoRA : Low-Load Sensitivity-Based Fine-Tuning for Large Language Models</i> Hao Zhang, Bo Huang, Zhenjia Li, Xi Xiao, Hui Yi Leong, Zumeng Zhang, Xinwei Long, Tianyang Wang and Hao Xu	13185
<i>ROSE: A Reward-Oriented Data Selection Framework for LLM Task-Specific Instruction Tuning</i> Yang Wu, Huayi Zhang, Yizheng Jiao, Lin Ma, Xiaozhong Liu, Jinhong Yu, Dongyu Zhang, Dezhi YU and Wei Xu.....	13200
<i>SimBA: Simplifying Benchmark Analysis Using Performance Matrices Alone</i> Nishant Subramani, Alfredo Gomez and Mona T. Diab	13220
<i>MarathiEmoExplain: A Dataset for Sentiment, Emotion, and Explanation in Low-Resource Marathi</i> Anuj Kumar, Mohammed Faisal Sayed, Satyadev Ahlawat and Yamuna Prasad.....	13234
<i>Active Domain Knowledge Acquisition with 100-Dollar Budget: Enhancing LLMs via Cost-Efficient, Expert-Involved Interaction in Sensitive Domains</i> Yang Wu, Raha Moraffah, Rujing Yao, Jinhong Yu, Zhimin Tao and Xiaozhong Liu	13244
<i>Structure-aware Propagation Generation with Large Language Models for Fake News Detection</i> Mengyang Chen, Lingwei Wei, Wei Zhou and Songlin Hu	13258
<i>UniCoM: A Universal Code-Switching Speech Generator</i> Sangmin Lee, Woojin Chung, Seyun Um and Hong-Goo Kang	13273
<i>Mitigating Sequential Dependencies: A Survey of Algorithms and Systems for Generation-Refinement Frameworks in Autoregressive Models</i> Yunhai Hu, Zining Liu, Zhenyuan Dong, Tianfan Peng, Bradley McDanel and Sai Qian Zhang	13289
<i>Do We Really Need All Those Dimensions? An Intrinsic Evaluation Framework for Compressed Embeddings</i> Nathan Inkiriwang, Necva Bölücü, Garth Tarr and Maciej Rybinski	13305
<i>Mixture of LoRA Experts for Continual Information Extraction with LLMs</i> Zitao Wang, Xinyi Wang and Wei Hu	13324
<i>Spelling-out is not Straightforward: LLMs' Capability of Tokenization from Token to Characters</i> Tatsuya Hiraoka and Kentaro Inui	13340
<i>OAgents: An Empirical Study of Building Effective Agents</i> He Zhu, Tianrui Qin, King Zhu, Heyuan Huang, Yeyi Guan, Jinxiang Xia, Hanhao Li, Yi Yao, Ningning Wang, Pai Liu, Tianhao Peng, Xin Gui, Li Xiaowan, Yuhui Liu, Xiangru Tang, Jian Yang, Ge Zhang, Xitong Gao, Yuchen Eleanor Jiang, Changwang Zhang, Jun Wang, Jiaheng Liu and Wangchunshu Zhou.....	13354
<i>2Columns1Row: A Russian Benchmark for Textual and Multimodal Table Understanding and Reasoning</i> Vildan Saburov, Daniil Vodolazsky, Danil Sazanakov and Alena Fenogenova	13370
<i>Permitted Knowledge Boundary: Evaluating the Knowledge-Constrained Responsiveness of Large Language Models</i> Wenrui Bao, Kai Wang, Siqiang Luo and Xiang Li	13390

<i>A Closer Look at Bias and Chain-of-Thought Faithfulness of Large (Vision) Language Models</i> Sriram Balasubramanian, Samyadeep Basu and Soheil Feizi	13406
<i>From Remembering to Metacognition: Do Existing Benchmarks Accurately Evaluate LLMs?</i> Geng Zhang, Yizhou Ying, Sihang Jiang, Jiaqing Liang, Guanglei Yue, Yifei Fu, Hailin Hu and Yanghua Xiao	13440
<i>How a Bilingual LM Becomes Bilingual: Tracing Internal Representations with Sparse Autoencoders</i> Tatsuro Inaba, Go Kamoda, Kentaro Inui, Masaru Isonuma, Yusuke Miyao, Yohei Oseki, Yu Takagi and Benjamin Heinzerling	13458
<i>MultiConIR: Towards Multi-Condition Information Retrieval</i> Xuan Lu, Sifan Liu, Bochao Yin, Yongqi Li, Xinghao Chen, Hui Su, Yaohui Jin, Wenjun Zeng and Xiaoyu Shen	13471
<i>HMCL: Task-Optimal Text Representation Adaptation through Hierarchical Contrastive Learning</i> Zhenyi Wang, Yapeng Jia, Haiyan Ning, Peng Wang, Dan Wang and Yitao Cao	13495
<i>KBAlign: Efficient Self Adaptation on Specific Textual Knowledge Bases</i> Zheni Zeng, Yuxuan Chen, Shi Yu, Ruobing Wang, Yukun Yan, Zhenghao Liu, Shuo Wang, Xu Han, Zhiyuan Liu and Maosong Sun	13519
<i>Revisiting Chain-of-Thought Prompting: Zero-shot Can Be Stronger than Few-shot</i> Xiang Cheng, Chengyan Pan, Minjun Zhao, Deyang Li, Fangchao Liu, Xinyu Zhang, Xiao Zhang and Yong Liu	13533
<i>RMTBench: Benchmarking LLMs Through Multi-Turn User-Centric Role-Playing</i> Hao Xiang, Tianyi Tang, Yang Su, Bowen Yu, An Yang, Fei Huang, Yichang Zhang, Yaojie Lu, Hongyu Lin, Xianpei Han, Jingren Zhou, Junyang Lin and Le Sun	13555
<i>Smart-Searcher: Incentivizing the Dynamic Knowledge Acquisition of LLMs via Reinforcement Learning</i> Huatong Song, Jinhao Jiang, Wenqing Tian, Zhipeng Chen, Yuhuan Wu, Jiahao Zhao, Yingqian Min, Xin Zhao, Lei Fang and Ji-Rong Wen	13572
<i>InteGround: On the Evaluation of Verification and Retrieval Planning in Integrative Grounding</i> Cheng Jiayang, Qianqian Zhuang, Haoran Li, Chunkit Chan, Xin Liu, Lin Qiu and Yangqiu Song 13587	
<i>MM-CRITIC: A Holistic Evaluation of Large Multimodal Models as Multimodal Critique</i> Gailun Zeng, Ziyang Luo, Hongzhan Lin, Yuchen Tian, Kaixin Li, Ziyang Gong, Jianxiong Guo and Jing Ma	13603
<i>On the Correspondence between the Squared Norm and Information Content in Text Embeddings</i> Enrique Amigo, Adrian Ghajari, Alejandro Benito-Santos and Diego De La Fuente Rodríguez 13631	
<i>Adversary-Aware DPO: Enhancing Safety Alignment in Vision Language Models via Adversarial Training</i> Fenghua Weng, Jian Lou, Jun Feng, Minlie Huang and Wenjie Wang	13644
<i>SLiNT: Structure-aware Language Model with Injection and Contrastive Training for Knowledge Graph Completion</i> Mengxue Yang, Chun Yang, Jiaqi Zhu, Jiafan Li, Jingqi Zhang, Yuyang Li and Ying Li . . .	13658

<i>LAVa: Layer-wise KV Cache Eviction with Dynamic Budget Allocation</i>	
Yiqun Shen, Song Yuan, Zhengze Zhang, Xiaoliang Wang, Daxin Jiang and Nguyen Cam-Tu	
13672	
<i>LoRA-PAR: A Flexible Dual-System LoRA Partitioning Approach to Efficient LLM Fine-Tuning</i>	
Yining Huang, Bin Li, Keke Tang and Meilian Chen	13693
<i>SimpleDeepSearcher: Deep Information Seeking via Web-Powered Reasoning Trajectory Synthesis</i>	
Shuang Sun, Huatong Song, Yuhao Wang, Ruiyang Ren, Jinhao Jiang, Junjie Zhang, Fei Bai, Jia Deng, Xin Zhao, Zheng Liu, Lei Fang, Zhongyuan Wang and Ji-Rong Wen	13705
<i>LLaVE: Large Language and Vision Embedding Models with Hardness-Weighted Contrastive Learning</i>	
Zhibin Lan, Liqiang Niu, Fandong Meng, Jie Zhou and Jinsong Su	13721
<i>SampleMix: A Sample-wise Pre-training Data Mixing Strategy by Coordinating Data Quality and Diversity</i>	
Xiangyu Xi, Deyang Kong, Jian Yang, Jiawei Yang, Zhengyu Chen, Wei Wang, Jingang Wang, Xunliang Cai, Shikun Zhang and Wei Ye	13736
<i>Evaluating Test-Time Scaling LLMs for Legal Reasoning: OpenAI o1, DeepSeek-R1, and Beyond</i>	
Yinghao Hu, Yaoyao Yu, Leilei Gan, Bin Wei, Kun Kuang and Fei Wu	13759
<i>LLM Agents for Education: Advances and Applications</i>	
Zhendong Chu, Shen Wang, Jian Xie, Tinghui Zhu, Yibo Yan, Jingheng Ye, Aoxiao Zhong, Xuming Hu, Jing Liang, Philip S. Yu and Qingsong Wen	13782
<i>Modeling Subjectivity in Cognitive Appraisal with Language Models</i>	
Yuxiang Zhou, Hainiu Xu, Desmond Ong, Maria Liakata, Petr Slovak and Yulan He	13811
<i>Dementia Through Different Eyes: Explainable Modeling of Human and LLM Perceptions for Early Awareness</i>	
Lotem Peled-Cohen, Maya Zadok, Nitay Calderon, Hila Gonen and Roi Reichart	13834
<i>Mitigating Hallucinations in Large Vision-Language Models by Self-Injecting Hallucinations</i>	
Yifan Lu, Ziqi Zhang, Chunfeng Yuan, Jun Gao, Congxuan Zhang, Xiaojuan Qi, Bing Li and Weiming Hu	13861
<i>How Much Do Large Language Models Know about Human Motion? A Case Study in 3D Avatar Control</i>	
Kunhang Li, Jason Naradowsky, Yansong Feng and Yusuke Miyao	13878
<i>The Search for Conflicts of Interest: Open Information Extraction in Scientific Publications</i>	
Garima Gaur, Oana Balalau, Ioana Manolescu and Prajna Upadhyay	13922
<i>On Collaborating Small and Large Models For Few-shot Intent Detection</i>	
Peng Chen and Bang Wang	13937
<i>A Survey on LLMs for Story Generation</i>	
Maria Teleki, Vedangi Bengali, Xiangjue Dong, Sai Tejas Janjur, Haoran Liu, Tian Liu, Cong Wang, Ting Liu, Yin Zhang, Frank Shipman and James Caverlee	13954
<i>From Knowledge to Treatment: Large Language Model Assisted Biomedical Concept Representation for Drug Repurposing</i>	
Chengrui Xiang, Tengfei Ma, Xiangzheng Fu, Yiping Liu, Bosheng Song and Xiangxiang Zeng	13967

<i>SKRAG: A Retrieval-Augmented Generation Framework Guided by Reasoning Skeletons over Knowledge Graphs</i>	
Xiaotong Xu, Yizhao Wang, Yunfei Liu and Shengyang Li	13983
<i>A Generative Framework for Personalized Sticker Retrieval</i>	
Changjiang Zhou, Ruqing Zhang, Jiafeng Guo, Yu-An Liu, Fan Zhang, Ganyuan Luo and Xueqi Cheng	13995
<i>Bridging Semantic and Modality Gaps in Zero-Shot Captioning via Retrieval from Synthetic Data</i>	
Zhiyue Liu and Wenkai Zhou	14010
<i>Humor in Pixels: Benchmarking Large Multimodal Models Understanding of Online Comics</i>	
Yuriel Ryan, Rui Yang Tan, Kenny Tsu Wei Choo and Roy Ka-Wei Lee	14024
<i>BiMediX2 : Bio-Medical EXpert LMM for Diverse Medical Modalities</i>	
Sahal Shaji Mullappilly, Mohammed Irfan Kurpath, Sara Pieri, Saeed Yahya Alseiari, Shanas Cholakal, Khaled M Aldahmani, Fahad Shahbaz Khan, Rao Muhammad Anwer, Salman Khan, Timothy Baldwin and Hisham Cholakal	14051
<i>DeMAC: Enhancing Multi-Agent Coordination with Dynamic DAG and Manager-Player Feedback</i>	
Yuhan Liu, Cong Xu, Lu Liu, Yihua Wang, Feiyu Chen, Qi Jia, Yaqian Zhao, Zhichun Wang and Xiang Li	14072
<i>Coherence of Argumentative Dialogue Snippets: A New Method for Large Scale Evaluation with an Application to Inference Anchoring Theory</i>	
Paul Piwek, Jacopo Amidei and Svetlana Stoyanchev	14099
<i>Angular Dispersion Accelerates k-Nearest Neighbors Machine Translation</i>	
Evgeniia Tokarchuk, Sergey Troshin and Vlad Niculae	14120
<i>Benchmarking Contextual and Paralinguistic Reasoning in Speech-LLMs: A Case Study with In-the-Wild Data</i>	
Qiongqiong Wang, Hardik Bhupendra Sailor, Tianchi Liu, Wenyu Zhang, Muhammad Huzaifah, Nattadaporn Lertcheva, Shuo Sun, Nancy F. Chen, Jinyang Wu and AiTi Aw	14133
<i>This is not a Disimprovement: Improving Negation Reasoning in Large Language Models via Prompt Engineering</i>	
Joshua Jose Dias Barreto and Abhik Jana	14149
<i>Make Every Letter Count: Building Dialect Variation Dictionaries from Monolingual Corpora</i>	
Robert Litschko, Verena Blaschke, Diana Burkhardt, Barbara Plank and Diego Frassinelli ..	14157
<i>SelfAug: Mitigating Catastrophic Forgetting in Retrieval-Augmented Generation via Distribution Self-Alignment</i>	
Yuqing Huang, Rongyang Zhang, Qimeng Wang, Chengqiang Lu, Yan Gao, Yiwu, Yao Hu, Xuyang Zhi, Guiquan Liu, Xin Li, Hao Wang and Enhong Chen	14175
<i>SEKE: Specialised Experts for Keyword Extraction</i>	
Matej Martinc, Thi Hong Hanh Tran, Senja Pollak and Boshko Koloski	14191
<i>$I+I>2$: A Synergistic Sparse and Low-Rank Compression Method for Large Language Models</i>	
Zeliang Zong, Kai Zhang, Zheyang Li, Wenming Tan, Ye Ren, Yiyan Zhai and Jilin Hu ...	14206
<i>InfIMM-WebMath-40B: Advancing Multimodal Pre-Training for Enhanced Mathematical Reasoning</i>	
Xiaotian Han, Yiren Jian, Xuefeng Hu, Haogeng Liu, Yiqi Wang, Qihang Fan, Yuang Ai, Huaibo Huang, Ran He, Zhenheng Yang and Quanzeng You	14221

<i>Zero-Shot Defense Against Toxic Images via Inherent Multimodal Alignment in LVLMs</i> Wei Zhao, Zhe Li, Yige Li and Jun Sun	14232
<i>Retrieval Augmented Generation based context discovery for ASR</i> Siskos Dimitrios, Stavros Papadopoulos, Pablo Peso Parada, Jisi Zhang, Karthikeyan Saravanan and Anastasios Drosou	14247
<i>pFedRAG: A Personalized Federated Retrieval-Augmented Generation System with Depth-Adaptive Tie- red Embedding Tuning</i> Hangyu He, Xin Yuan, Kai Wu, Ren Ping Liu and Wei Ni.....	14255
<i>ReCUT: Balancing Reasoning Length and Accuracy in LLMs via Stepwise Trails and Preference Opti- mization</i> Zhensheng Jin, Xinze Li, Yifan Ji, Chunyi Peng, Zhenghao Liu, Qi Shi, Yukun Yan, Shuo Wang, Furong Peng and Ge Yu	14269
<i>CURE: Controlled Unlearning for Robust Embeddings — Mitigating Conceptual Shortcuts in Pre- Trained Language Models</i> Aysenur Kocak, Shuo Yang, Bardh Prenkaj and Gjergji Kasneci	14283
<i>MLAlgo-Bench: Can Machines Implement Machine Learning Algorithms?</i> Yunfei Wang, Yeqin Zhang, Yuyang Wu, Liang Lu, Phi Le Nguyen, Xiaoliang Wang and Nguyen Cam-Tu.....	14298
<i>Fair Text-Attributed Graph Representation Learning</i> Ruilin Luo, Tianle Gu, Lin Wang, Yunfeng Zhou, Songtao Jiang, Lei Wang and Yujiu Yang.....	14330
<i>Human-Inspired Obfuscation for Model Unlearning: Local and Global Strategies with Hyperbolic Re- presentations</i> Zekun Wang, Jingjie Zeng, Yingxu Li, Liang Yang and Hongfei Lin	14354
<i>Do Influence Functions Work on Large Language Models?</i> Zhe Li, Wei Zhao, Yige Li and Jun Sun	14367
<i>TRUEBench: Can LLM Response Meet Real-world Constraints as Productivity Assistant?</i> Jiho Park, Jongyoon Song, Minjin Choi, Kyuho Heo, Taehun Huh and Ji Won Kim	14383
<i>CausalMACE: Causality Empowered Multi-Agents in Minecraft Cooperative Tasks</i> Qi Chai, Zhang Zheng, Junlong Ren, Deheng Ye, Zichuan Lin and Hao Wang.....	14410
<i>Harry Potter is Still Here! Probing Knowledge Leakage in Targeted Unlearned Large Language Models</i> Bang Trinh Tran To and Thai Le.....	14427
<i>Learning Trajectories of Figurative Language for Pre-Trained Language Models</i> Nicola Arici, Luca Putelli, Ejdis Gjinika, Ivan Serina and Alfonso Gerevini	14440
<i>BcQLM: Efficient Vision-Language Understanding with Distilled Q-Gated Cross-Modal Fusion</i> Sike Xiang, Shuang Chen and Amir Atapour-Abarghouei	14462
<i>HapticCap: A Multimodal Dataset and Task for Understanding User Experience of Vibration Haptic Signals</i> Guimin Hu, Daniel Hershcovich and Hasti Seifi.....	14473
<i>SubDocTrans: Enhancing Document-level Machine Translation with Plug-and-play Multi-granularity Knowledge Augmentation</i> Hanghai Hong, Yibo Xie, Jiawei Zheng and Xiaoli Wang	14490

<i>Social Bias Evaluation for Large Language Models Requires Prompt Variations</i>	
Rem Hida, Masahiro Kaneko and Naoaki Okazaki	14507
<i>Training with Fewer Bits: Unlocking Edge LLMs Training with Stochastic Rounding</i>	
Taowen Liu, Marta Andronic, Deniz Gunduz and George Anthony Constantinides	14531
<i>FactReasoner: A Probabilistic Approach to Long-Form Factuality Assessment for Large Language Models</i>	
Radu Marinescu, Debarun Bhattacharjya, Junkyu Lee, Tigran T. Tehrakian, Javier Carnerero-Cano, Yufang Hou, Elizabeth M. Daly and Alessandra Pascale	14547
<i>Robust Knowledge Editing via Explicit Reasoning Chains for Distractor-Resilient Multi-Hop QA</i>	
Yuchen Wu, Liang Ding, Li Shen and Dacheng Tao	14578
<i>RadialRouter: Structured Representation for Efficient and Robust Large Language Models Routing</i>	
Ruihan Jin, Pengpeng Shao, Zhengqi Wen, Jinyang Wu, Mingkuan Feng, Shuai Zhang and Jianhua Tao	14587
<i>Decoding Uncertainty: The Impact of Decoding Strategies for Uncertainty Estimation in Large Language Models</i>	
Wataru Hashimoto, Hidetaka Kamigaito and Taro Watanabe	14601
<i>Elucidating Mechanisms of Demographic Bias in LLMs for Healthcare</i>	
Hiba Ahsan, Arnab Sen Sharma, Silvio Amir, David Bau and Byron C Wallace	14614
<i>Can You Trick the Grader? Adversarial Persuasion of LLM Judges</i>	
Yerin Hwang, Dongryeol Lee, Taegwan Kang, Yongil Kim and Kyomin Jung	14632
<i>Navigating the Unknown: Intent Classification and Out-of-Distribution Detection Using Large Language Models</i>	
Yusuf Sali and Sitki Can Toraman	14652
<i>Trust Me, I'm Wrong: LLMs Hallucinate with Certainty Despite Knowing the Answer</i>	
Adi Simhi, Itay Itzhak, Fazl Barez, Gabriel Stanovsky and Yonatan Belinkov	14665
<i>QUARTZ: QA-based Unsupervised Abstractive Refinement for Task-oriented Dialogue Summarization</i>	
Mohamed Imed Eddine Ghebriout, Gaël Guibon, Ivan Lerner and Emmanuel Vincent	14689
<i>MDSEval: A Meta-Evaluation Benchmark for Multimodal Dialogue Summarization</i>	
Yinhong Liu, Jianfeng He, Hang Su, Ruixue Lian, Yi Nian, Jake W. Vincent, Srikanth Vishnubhotla, Robinson Piramuthu and Saab Mansour	14707
<i>PMPO: Probabilistic Metric Prompt Optimization for Small and Large Language Models</i>	
ChenZhuo Zhao, Ziqian Liu, Xinda Wang, Junting Lu and Chaoyi Ruan	14728
<i>Evaluating the Creativity of LLMs in Persian Literary Text Generation</i>	
Armin Tourajmehr, Mohammad Reza Modarres and Yadollah Yaghoobzadeh	14762
<i>SCDTour: Embedding Axis Ordering and Merging for Interpretable Semantic Change Detection</i>	
Taichi Aida and Danushka Bollegala	14775
<i>Resolving UnderEdit & OverEdit with Iterative & Neighbor-Assisted Model Editing</i>	
Bhiman Kumar Baghel, Emma Jordan, Zheyuan Ryan Shi and Xiang Lorraine Li	14786
<i>LLM-empowered Dynamic Prompt Routing for Vision-Language Models Tuning under Long-Tailed Distributions</i>	
Yongju Jia, Jiarui Ma, Xiangxian Li, Baiqiao Zhang, Xianhui Cao, Juan Liu and Yulong Bian	14809

<i>HGAdapter: Hypergraph-based Adapters in Language Models for Code Summarization and Clone Detection</i>	
Guang Yang and Yujie Zhu	14823
<i>Evaluating distillation methods for data-efficient syntax learning</i>	
Takateru Yamakoshi, Thomas L. Griffiths, R. Thomas McCoy and Robert D. Hawkins	14834
<i>Going to a trap house conveys more fear than Going to a mall": Benchmarking Emotion Context Sensitivity for LLMs</i>	
Eojin Jeon, Mingyu Lee, Sangyun Kim, Junho Kim, Wansee Cho, Tae-Eui Kam and SangKeun Lee	14848
<i>[MASK]ED - Language Modeling for Explainable Classification and Disentangling of Socially Unacceptable Discourse.</i>	
Dimitra Niaouri, Mohamed Rayane Ghilene, Michele Linardi and Julien Longhi	14870
<i>A Survey of Cognitive Distortion Detection and Classification in NLP</i>	
Archie Sage, Jeroen Keppens and Helen Yannakoudakis	14884
<i>Curse of Knowledge: Your Guidance and Provided Knowledge are biasing LLM Judges in Complex Evaluation</i>	
Weiyuan Li, Xintao Wang, Siyu Yuan, Rui Xu, Jiangjie Chen, Qingqing Dong, Yanghua Xiao and Deqing Yang	14900
<i>Self-Training Large Language Models with Confident Reasoning</i>	
Hyosoon Jang, Yunhui Jang, Sungjae Lee, Jungseul Ok and Sungsoo Ahn	14925
<i>Error Typing for Smarter Rewards: Improving Process Reward Models with Error-Aware Hierarchical Supervision</i>	
Tej Deep Pala, Panshul Sharma, Amir Zadeh, Chuan Li and Soujanya Poria	14940
<i>Enhancing LLM-Based Persuasion Simulations with Cultural and Speaker-Specific Information</i>	
Weicheng Ma, Hefan Zhang, Shiyu Ji, Farnoosh Hashemi, Qichao Wang, Ivory Yang, Joice Chen, Juanwen Pan, Michael Macy, Saeed Hassanpour and Soroush Vosoughi	14955
<i>An LLM-based Temporal-spatial Data Generation and Fusion Approach for Early Detection of Late Onset Alzheimer's Disease (LOAD) Stagings Especially in Chinese and English-speaking Populations</i>	
Yang Han, Jacqueline C.k. Lam, Victor O.k. Li and Lawrence Y. L. Cheung	14977
<i>Side Effects of Erasing Concepts from Diffusion Models</i>	
Shaswati Saha, Sourajit Saha, Manas Gaur and Tejas Gokhale	14991
<i>SaCa: A Highly Compatible Reinforcing Framework for Knowledge Graph Embedding via Structural Pattern Contrast</i>	
Jiashi Lin, Changhong Jiang, Yixiao Wang, Xinyi Zhu, Zhongtian Hu and Wei Zhang	15008
<i>Real, Fake, or Manipulated? Detecting Machine-Influenced Text</i>	
Yitong Wang, Zhongping Zhang, Margherita Piana, Zheng Zhou, Peter Gerstoft and Bryan A. Plummer	15022
<i>Character is Destiny: Can Persona-assigned Language Models Make Personal Choices?</i>	
Rui Xu, Xintao Wang, Jiangjie Chen, Siyu Yuan, Xinfeng Yuan, Jiaqing Liang, Zulong Chen, Xiaoqingdong and Yanghua Xiao	15038
<i>Neutral Is Not Unbiased: Evaluating Implicit and Intersectional Identity Bias in LLMs Through Structured Narrative Scenarios</i>	
Saba Ghanbari Haez and Mauro Dragoni	15060

<i>BTW: A Non-Parametric Variance Stabilization Framework for Multimodal Model Integration</i> Jun Hou, Le Wang and Xuan Wang	15089
<i>Can LLMs Be Efficient Predictors of Conversational Derailment?</i> Kaustubh Olpadkar, Vikram Sunil Bajaj and Leslie Barrett	15104
<i>Q-PRM: Adaptive Query Rewriting for Retrieval-Augmented Generation via Step-level Process Supervision</i> Xiaopeng Ye, Chen Xu, Chaoliang Zhang, Zhaocheng Du, Jun Xu, Gang Wang and Zhenhua Dong	15113
<i>Factuality Beyond Coherence: Evaluating LLM Watermarking Methods for Medical Texts</i> Rochana Prih Hastuti, Rian Adam Rajagede, Mansour Al Ghanim, Mengxin Zheng and Qian Lou	15129
<i>Guess What I am Thinking: A Benchmark for Inner Thought Reasoning of Role-Playing Language Agents</i> Rui Xu, Mingyu Wang, Xintao Wang, Dakuan Lu, Xiaoyu Tan, Wei Chu and Xu Yinghui .	15148
<i>Dropping Experts, Recombining Neurons: Retraining-Free Pruning for Sparse Mixture-of-Experts LLMs</i> Yixiao Zhou, Ziyu Zhao, Dongzhou Cheng, Zhiliang wu, Jie Gui, Yi Yang, Fei Wu, Yu Cheng and Hehe Fan	15169
<i>BiasFilter: An Inference-Time Debiasing Framework for Large Language Models</i> Xiaoqing Cheng, Ruizhe Chen, Hongying Zan, Yuxiang Jia and Min Peng	15187
<i>X-LeBench: A Benchmark for Extremely Long Egocentric Video Understanding</i> Wenqi Zhou, Kai Cao, Hao Zheng, Yunze Liu, Xinyi Zheng, Miao Liu, Per Ola Kristensson, Walterio W. Mayol-Cuevas, Fan Zhang, Weizhe Lin and Junxiao Shen	15206
<i>A Survey on Multi-modal Intent Recognition: Recent Advances and New Frontiers</i> Zhihong Zhu, Fan Zhang, Yunyan Zhang, Jinghan Sun, Zhiqi Huang, Qingqing Long, Bowen Xing and Xian Wu	15223
<i>Will Annotators Disagree? Identifying Subjectivity in Value-Laden Arguments</i> Amir Homayounirad, Enrico Liscio, Tong Wang, Catholijn M Jonker and Luciano Cavalcante Siebert	15237
<i>LLMs Can Compensate for Deficiencies in Visual Representations</i> Sho Takishita, Jay Gala, Abdelrahman Mohamed, Kentaro Inui and Yova Kementchedjheva	15253
<i>Adapting Large Language Models for Character-based Augmentative and Alternative Communication</i> Dylan Gaines and Keith Vertanen	15273
<i>Token-Level Metrics for Detecting Incorrect Gold Annotations in Named Entity Recognition</i> Elena Merdjanovska and Alan Akbik	15292
<i>Exploring Paraphrasing Strategies for CEFR A1-Level Constraints in LLMs</i> Eugenio Marzona, Maria Goikhman, Alessio Palmero Aprosio and Massimo Zancanaro ..	15305
<i>Efficient Layer-wise LLM Fine-tuning for Revision Intention Prediction</i> Zhexiong Liu and Diane Litman	15319
<i>ConText-LE: Cross-Distribution Generalization for Longitudinal Experiential Data via Narrative-Based LLM Representations</i> Ahatsham Hayat, Bilal Khan and Mohammad Rashedul Hasan	15335

<i>Chain of Strategy Optimization Makes Large Language Models Better Emotional Supporter</i> Weixiang Zhao, Xingyu Sui, Xinyang Han, Yang Deng, Yulin Hu, Jiahe Guo, Libo Qin, Qianyun Du, Shijin Wang, Yanyan Zhao, Bing Qin and Ting Liu	15361
<i>Unlocking Legal Knowledge: A Multilingual Dataset for Judicial Summarization in Switzerland</i> Luca Rolshoven, Vishvakshenan Rasiah, Srinanda Brügger Bose, Sarah Hostettler, Lara Burkhalter, Matthias Stürmer and Joel Niklaus	15382
<i>Context Minimization for Resource-Constrained Text Classification: Optimizing Performance-Efficiency Trade-offs through Linguistic Features</i> Nahid Hossain and Md Faisal Kabir	15412
<i>FLAIRR-TS - Forecasting LLM-Agents with Iterative Refinement and Retrieval for Time Series</i> Gunjan Jalori, Preetika Verma and Sercan O Arik	15427
<i>ULTRABENCH: Benchmarking LLMs under Extreme Fine-grained Text Generation</i> Longfei Yun, Letian Peng and Jingbo Shang	15438
<i>The Price of Format: Diversity Collapse in LLMs</i> Longfei Yun, Chenyang An, Zilong Wang, Letian Peng and Jingbo Shang	15454
<i>Zipf's and Heaps' Laws for Tokens and LLM-generated Texts</i> Nikolay Mikhaylovskiy	15469
<i>LLMs for Bayesian Optimization in Scientific Domains: Are We There Yet?</i> Rushil Gupta, Jason Hartford and Bang Liu	15482
<i>A Comprehensive Taxonomy of Negation for NLP and Neural Retrievers</i> Roxana Petcu, Samarth Bhargav, Maarten de Rijke and Evangelos Kanoulas	15511
<i>Identifying Noise in Human-Created Datasets using Training Dynamics from Generative Models</i> Maeda Hanafi, Ishan Jindal, Yannis Katsis, Lucian Popa and Huaiyu Zhu	15534
<i>Can Multiple Responses from an LLM Reveal the Sources of Its Uncertainty?</i> Yang Nan, Pengfei He, Ravi Tandon and Han Xu	15551
<i>AfroXLMR-Social: Adapting Pre-trained Language Models for African Languages Social Media Text</i> Tadesse Destaw Belay, Israel Abebe Azime, Ibrahim Said Ahmad, David Ifeoluwa Adelani, Idris Abdulmumin, Abinew Ali Ayele, Shamsuddeen Hassan Muhammad and Seid Muhie Yimam ...	15570
<i>Teaching Language Models To Gather Information Proactively</i> Tenghao Huang, Sihao Chen, Muhao Chen, Jonathan May, Longqi Yang, Mengting Wan and Pei Zhou	15588
<i>Linguistic Alignment Predicts Learning in Small Group Tutoring Sessions</i> Dorothea French, Robert Moulder, Kelechi Ezema, Katharina von der Wense and Sidney K. D'Mello	15600
<i>EfficientXLang: Towards Improving Token Efficiency Through Cross-Lingual Reasoning</i> Sanchit Ahuja, Praneetha Vaddamanu and Barun Patra	15612
<i>Not Lost After All: How Cross-Encoder Attribution Challenges Position Bias Assumptions in LLM Summarization</i> Elahe Rahimi, Hassan Sajjad, Domenic Rosati, Abeer Badawi, Elham Dolatabadi and Frank Rudzicz	15625
<i>FuzzAug: Data Augmentation by Coverage-guided Fuzzing for Neural Test Generation</i> Yifeng He, Jicheng Wang, Yuyang Rong and Hao Chen	15642

<i>DrAgent: Empowering Large Language Models as Medical Agents for Multi-hop Medical Reasoning</i> Fenglin Liu, Zheng Li, Hongjian Zhou, Qingyu Yin, Jingfeng Yang, Xin Liu, Zhengyang Wang, Xianfeng Tang, Shiyang Li, Xiang He, Ruijie Wang, Bing Yin, Xiao Gu, Lei Clifton and David A. Clifton	15656
<i>XRAG: Cross-lingual Retrieval-Augmented Generation</i> Wei Liu, Sony Trenous, Leonardo F. R. Ribeiro, Bill Byrne and Felix Hieber	15669
<i>Can VLMs Recall Factual Associations From Visual References?</i> Dhananjay Ashok, Ashutosh Chaubey, Hirona Jacqueline Arai, Jonathan May and Jesse Thoma- son	15691
<i>MFTCXplain: A Multilingual Benchmark Dataset for Evaluating the Moral Reasoning of LLMs through Multi-hop Hate Speech Explanation</i> Jackson Trager, Francielle Vargas, Diego Alves, Matteo Guida, Mikel K. Ngueajio, Ameeta Agrawal, Yalda Daryani, Farzan Karimi Malekabadi and Flor Miriam Plaza-del-Arco	15709
<i>Large Language Models for Multilingual Previously Fact-Checked Claim Detection</i> Ivan Vykopal, Matúš Pikuliak, Simon Ostermann, Tatiana Anikina, Michal Gregor and Marian Simko	15741
<i>Debating for Better Reasoning in Vision-Language Models</i> Ashutosh Adhikari and Mirella Lapata	15766
<i>Fine-tuning LLMs with Cross-Attention-based Weight Decay for Bias Mitigation</i> Farsheed Haque, Zhe Fu, Depeng Xu, Shuhan Yuan and Xi Niu	15785
<i>Profiling LLM's Copyright Infringement Risks under Adversarial Persuasive Prompting</i> Jikai Long, Ming Liu, Xiusi Chen, Jialiang Xu, Shenglan Li, Zhaozhuo Xu and Denghui Zhang 15799	
<i>Residualized Similarity for Faithfully Explainable Authorship Verification</i> Peter Zeng, Pegah Alipoormolabashi, Jihu Mun, Gourab Dey, Nikita Soni, Niranjan Balasubra- manian, Owen Rambow and H. Schwartz	15824
<i>Post-hoc Study of Climate Microtargeting on Social Media Ads with LLMs: Thematic Insights and Fairness Evaluation</i> Tunazzina Islam and Dan Goldwasser	15838
<i>MRFD: Multi-Region Fusion Decoding with Self-Consistency for Mitigating Hallucinations in LLMs</i> Haonan Ge, Yiwei Wang, Ming-Hsuan Yang and Yujun Cai	15860
<i>SIMBA UQ: Similarity-Based Aggregation for Uncertainty Quantification in Large Language Models</i> Debarun Bhattacharjya, Balaji Ganesan, Junkyu Lee, Radu Marinescu, Katya Mirylenka, Michael Glass and Xiao Shou	15880
<i>Mind the Dialect: NLP Advancements Uncover Fairness Disparities for Arabic Users in Recommenda- tion Systems</i> Abdulla Alshabanah and Murali Annavaram	15895
<i>Hopscotch: Discovering and Skipping Redundancies in Language Models</i> Mustafa Eyceoz, Nikhil Shivakumar Nayak, Hao Wang, Ligong Han and Akash Srivastava	15904
<i>CLEAR: A Clinically Grounded Tabular Framework for Radiology Report Evaluation</i> Yuyang Jiang, Chacha Chen, Shengyuan Wang, Feng Li, Zecong Tang, Benjamin M. Mervak, Lydia Chelala, Christopher M Straus, Reve Chahine, Samuel G. Armato Iii and Chenhao Tan ...	15914

<i>Parsing the Switch: LLM-Based UD Annotation for Complex Code-Switched and Low-Resource Languages</i>	
Olga Kellert, Nemika Tyagi, Muhammad Imran, Nelvin Licon-Guevara and Carlos Gómez-Rodríguez	15934
<i>HetGCoT: Heterogeneous Graph-Enhanced Chain-of-Thought LLM Reasoning for Academic Question Answering</i>	
Runsong Jia, Mengjia Wu, Ying Ding, Jie Lu and Yi Zhang	15950
<i>S*: Test Time Scaling for Code Generation</i>	
Dacheng Li, Shiyi Cao, Chengkun Cao, Xiuyu Li, Shangyin Tan, Kurt Keutzer, Jiarong Xing, Joseph E. Gonzalez and Ion Stoica	15964
<i>Language Models Can Easily Learn to Reason from Demonstrations</i>	
Dacheng Li, Shiyi Cao, Tyler Griggs, Shu Liu, Xiangxi Mo, Eric Tang, Sumanth Hegde, Kourosh Hakhamaneshi, Shishir G Patil, Matei Zaharia, Joseph E. Gonzalez and Ion Stoica	15979
<i>FSTs vs ICL: Generalisation in LLMs for an under-resourced language</i>	
Ximena Gutierrez, Mikel Segura Elizalde and Victor Mijangos	15998
<i>SRM-LLM: Semantic Relationship Mining with LLMs for Temporal Knowledge Graph Extrapolation</i>	
Fu Zhang, Panfeng Zhang and Jingwei Cheng	16007
<i>Captioning for Text-Video Retrieval via Dual-Group Direct Preference Optimization</i>	
Ji Soo Lee, Byungoh Ko, Jaewon Cho, Howoong Lee, Jaewoon Byun and Hyunwoo J. Kim	16022
<i>Benchmarking and Improving LLM Robustness for Personalized Generation</i>	
Chimaobi Okite, Naihao Deng, Kiran Bodipati, Huaidian Hou, Joyce Chai and Rada Mihalcea	16040
<i>MemeInterpret: Towards an All-in-One Dataset for Meme Understanding</i>	
Jeongsik Park, Khoi P. N. Nguyen, Jihyung Park, Minseok Kim, Jaeheon Lee, Jae Won Choi, Kalyani Ganta, Phalgun Ashrit Kasu, Rohan Sarakinti, Sanjana Vippera, Sai Sathanapalli, Nishan Vaghani and Vincent Ng	16073
<i>CoRAG: Enhancing Hybrid Retrieval-Augmented Generation through a Cooperative Retriever Architecture</i>	
Zaiyi Zheng, Song Wang, Zihan Chen, Yaochen Zhu, Yinhan He, Liangjie Hong, Qi Guo and Jundong Li	16088
<i>Hallucination Detection in Structured Query Generation via LLM Self-Debating</i>	
Miaoran Li, Jiangning Chen, Minghua Xu and Xiaolong Wang	16102
<i>Not All Options Are Created Equal: Textual Option Weighting for Token-Efficient LLM-Based Knowledge Tracing</i>	
Jongwoo Kim, SeongYeub Chu, Bryan Wong and Mun Yong Yi	16114
<i>Public Data Assisted Differentially Private In-Context Learning</i>	
Seongho Joo, Hyukhun Koh and Kyomin Jung	16129
<i>Inducing Argument Facets for Faithful Opinion Summarization</i>	
Jian Wang, Yanjie Liang, Yuqing Sun and Bin Gong	16153
<i>Scaling Laws Are Unreliable for Downstream Tasks: A Reality Check</i>	
Nicholas Lourie, Michael Y. Hu and Kyunghyun Cho	16167

<i>Familiarity-Aware Evidence Compression for Retrieval-Augmented Generation</i>	Dongwon Jung, Qin Liu, Tenghao Huang, Ben Zhou and Muhao Chen	16181
<i>O_O-VC: Synthetic Data-Driven One-to-One Alignment for Any-to-Any Voice Conversion</i>	Huu Tuong Tu, Huan Vu, Cuong Tien Nguyen, Dien Hy Ngo and Nguyen Thi Thu Trang .	16197
<i>Simple Factuality Probes Detect Hallucinations in Long-Form Natural Language Generation</i>	Jiatong Han, Neil Band, Muhammed Razzak, Jannik Kossen, Tim G. J. Rudner and Yarin Gal	16209
<i>CESRec: Constructing Pseudo Interactions for Sequential Recommendation via Conversational Feedback</i>	Yifan Wang, Shen Gao, Jiabao Fang, Rui Yan, Billy Chiu and Shuo Shang	16227
<i>TTPA: Token-level Tool-use Preference Alignment Training Framework with Fine-grained Evaluation</i>	Chengrui Huang, Shen Gao, Zhengliang Shi, Dongsheng Wang and Shuo Shang	16240
<i>Avoiding Knowledge Edit Skipping in Multi-hop Question Answering with Guided Decomposition</i>	Yi Liu, Xiangrong Zhu, Xiangyu Liu, Wei Wei and Wei Hu	16256
<i>Bridging the Creativity Understanding Gap: Small-Scale Human Alignment Enables Expert-Level Humor Ranking in LLMs</i>	Kuan Lok Zhou, Jiayi Chen, Siddharth Suresh, Reuben Narad, Timothy T. Rogers, Lalit K Jain, Robert D Nowak, Bob Mankoff and Jifan Zhang	16273
<i>SMARTMiner: Extracting and Evaluating SMART Goals from Low-Resource Health Coaching Notes</i>	Iva Bojic, Qi Chwen Ong, Stephanie Hilary Xinyi Ma, Lin Ai, Zheng Liu, Ziwei Gong, Julia Hirschberg, Andy Hau Yan HO and Andy W. H. Khong	16288
<i>GRIL: Knowledge Graph Retrieval-Integrated Learning with Large Language Models</i>	Jialin Chen, Houyu Zhang, Seongjun Yun, Alejandro Mottini, Rex Ying, Xiang Song, Vassilis N. Ioannidis, Zheng Li and Qingjun Cui	16306
<i>Exploring Deductive and Inductive Reasoning Capabilities of Large Language Models in Procedural Planning</i>	Jiabao Kang, Xinye Li, Liyan Xu, Qingbin Liu, Xi Chen, Zhiying Tu, Dianhui Chu and Dianbo Sui	16320
<i>KELE: A Multi-Agent Framework for Structured Socratic Teaching with Large Language Models</i>	Xian Peng, Pan Yuan, Dong Li, Junlong Cheng, Qin Fang and Zhi Liu	16342
<i>VisualEDU: A Benchmark for Assessing Coding and Visual Comprehension through Educational Problem-Solving Video Generation</i>	Hao Chen, Tianyu Shi, Pengran Huang, Zeyuan Li, Jiahui Pan, Qianglong Chen and Lewei He	16363
<i>OkraLong: A Flexible Retrieval-Augmented Framework for Long-Text Question Answering</i>	Yulong Hui, Yihao Liu, Yao Lu and Huanchen Zhang	16395
<i>VerifiAgent: a Unified Verification Agent in Language Model Reasoning</i>	Jiuzhou Han, Wray Buntine and Ehsan Shareghi	16410
<i>DrKGC: Dynamic Subgraph Retrieval-Augmented LLMs for Knowledge Graph Completion across General and Biomedical Domains</i>	Yongkang Xiao, Sinian Zhang, Yi Dai, Huixue Zhou, Jue Hou, Jie Ding and Rui Zhang . . .	16432

<i>Understanding the Language Model to Solve the Symbolic Multi-Step Reasoning Problem from the Perspective of Buffer Mechanism</i>	
Zhiwei Wang, Yunji Wang, Zhongwang Zhang, Zhangchen Zhou, Hui Jin, Tianyang Hu, Jiacheng Sun, Zhenguo Li, Yaoyu Zhang and Zhi-Qin John Xu	16446
<i>TwT: Thinking without Tokens by Habitual Reasoning Distillation with Multi-Teachers' Guidance</i>	
Jingxian Xu, Mengyu Zhou, Weichang Liu, Hanbing Liu, Shi Han and Dongmei Zhang . . .	16475
<i>DAVIS: Planning Agent with Knowledge Graph-Powered Inner Monologue</i>	
Minh Pham Dinh, Michael G Yankoski, Munira Syed and Trenton W. Ford	16490
<i>When Instructions Multiply: Measuring and Estimating LLM Capabilities of Multiple Instructions Following</i>	
Keno Harada, Yudai Yamazaki, Masachika Taniguchi, Edison Marrese-Taylor, Takeshi Kojima, Yusuke Iwasawa and Yutaka Matsuo	16506
<i>FormosanBench: Benchmarking Low-Resource Austronesian Languages in the Era of Large Language Models</i>	
Kaiying Kevin Lin, Hsi-Yu Chen and Haopeng Zhang	16527
<i>SeaPO: Strategic Error Amplification for Robust Preference Optimization of Large Language Models</i>	
Jun Rao, Yunjie Liao, Xuebo Liu, Zepeng Lin, Lian Lian, Dong Jin, Shengjun Cheng, Jun Yu and Min Zhang	16540
<i>FigEx: Aligned Extraction of Scientific Figures and Captions</i>	
Jifeng Song, Arun Das, Ge Cui and Yufei Huang	16558
<i>PATIMT-Bench: A Multi-Scenario Benchmark for Position-Aware Text Image Machine Translation in Large Vision-Language Models</i>	
Wanru Zhuang, Wenbo Li, Zhibin Lan, Xu Han, Peng Li and Jinsong Su	16572
<i>Safeguard Fine-Tuned LLMs Through Pre- and Post-Tuning Model Merging</i>	
Hua Farn, Hsuan Su, Shachi H. Kumar, Saurav Sahay, Shang-Tse Chen and Hung-yi Lee .	16589
<i>Self-Ensemble: Mitigating Confidence Distortion for Large Language Models</i>	
Zicheng Xu, Guanchu Wang, Guangyao Zheng, Yu-Neng Chuang, Alex Szalay, Xia Hu and Vladimir Braverman	16603
<i>Annotation-Efficient Language Model Alignment via Diverse and Representative Response Texts</i>	
Yuu Jinnai and Ukyo Honda	16616
<i>Explainable Chain-of-Thought Reasoning: An Empirical Analysis on State-Aware Reasoning Dynamics</i>	
Sheldon Yu, Yuxin Xiong, Junda Wu, Xintong Li, Tong Yu, Xiang Chen, Ritwik Sinha, Jingbo Shang and Julian McAuley	16660
<i>DecisionFlow: Advancing Large Language Model as Principled Decision Maker</i>	
Xiusi Chen, Shanyong Wang, Cheng Qian, Hongru Wang, Peixuan Han and Heng Ji	16668
<i>M-Ped: Multi-Prompt Ensemble Decoding for Large Language Models</i>	
Jiaxin Guo, Daimeng Wei, Yuanchang Luo, Hengchao Shang, Zongyao Li, Jinlong Yang, Zhanglin Wu, Zhiqiang Rao, Shimin Tao and Hao Yang	16693
<i>Butterfly Effects in Toolchains: A Comprehensive Analysis of Failed Parameter Filling in LLM Tool-Agent Systems</i>	
Qian Xiong, Yuekai Huang, Ziyu Jiang, Zhiyuan Chang, Yujia Zheng, Tianhao Li and Mingyang Li	16712

<i>FinLFQA: Evaluating Attributed Text Generation of LLMs in Financial Long-Form Question Answering</i> Yitao Long, Tiansheng Hu, Yilun Zhao, Arman Cohan and Chen Zhao	16730
<i>BenchMAX: A Comprehensive Multilingual Evaluation Suite for Large Language Models</i> Xu Huang, Wenhao Zhu, Hanxu Hu, Conghui He, Lei Li, Shujian Huang and Fei Yuan . . .	16751
<i>Assessing the Sensitivity and Alignment of FOL Closeness Metrics</i> Ramya Keerthy Thatikonda, Wray Buntine and Ehsan Shareghi	16775
<i>FoodSafeSum: Enabling Natural Language Processing Applications for Food Safety Document Summarization and Analysis</i> Juli Bakagianni, Korbinian Randl, Guido Rocchietti, Cosimo Rulli, Franco Maria Nardini, Salvatore Trani, Aron Henriksson, Anna Romanova and John Pavlopoulos	16786
<i>Self-adaptive Dataset Construction for Real-World Multimodal Safety Scenarios</i> Jingen Qu, Lijun Li, Bo Zhang, Yichen Yan and Jing Shao	16805
<i>EnDive: A Cross-Dialect Benchmark for Fairness and Performance in Large Language Models</i> Abhay Gupta, Jacob Cheung, Philip Meng, Shayan Sayyed, Kevin Zhu, Austen Liao and Sean O'Brien	16830
<i>FAEDKV: Infinite-Window Fourier Transform for Unbiased KV Cache Compression</i> Runchao Li, Yao Fu, Mu Sheng, Xianxuan Long, Haotian Yu and Pan Li	16856
<i>Dynamic Injection of Entity Knowledge into Dense Retrievers</i> Ikuya Yamada, Ryokan Ri, Takeshi Kojima, Yusuke Iwasawa and Yutaka Matsuo	16867
<i>When Personalization Meets Reality: A Multi-Faceted Analysis of Personalized Preference Learning</i> Yijiang River Dong, Tiancheng Hu, Yinhong Liu, Ahmet Üstün and Nigel Collier	16880
<i>MASTER: Multi-Agent Security Through Exploration of Roles and Topological Structures - A Comprehensive Framework</i> Yifan Zhu, Chao Zhang, Xin Shi, Xueqiao Zhang, Yi Yang and Yawei Luo	16895
<i>MONAQ: Multi-Objective Neural Architecture Querying for Time-Series Analysis on Resource-Constrained Devices</i> Patara Trirat and Jae-Gil Lee	16922
<i>StandUp4AI: A New Multilingual Dataset for Humor Detection in Stand-up Comedy Videos</i> Valentin Barriere, Nahuel Gomez, Léo Hemamou, Sofia Callejas and Brian Ravenet	16951
<i>Does Visual Grounding Enhance the Understanding of Embodied Knowledge in Large Language Models?</i> Zhihui Yang, Yupei Wang, Kaijie Mo, Zhe Zhao and Renfen Hu	16960
<i>Semantic Contribution-Aware Adaptive Retrieval for Black-Box Models</i> Qinhong Lin, Zhongliang Yang, Yuang Cai, Dingfu Yu, Xuan Xu, Yu Li and Linna Zhou .	16979
<i>On Guardrail Models' Robustness to Mutations and Adversarial Attacks</i> Elias Bassani and Ignacio Sanchez	16995
<i>IP-Dialog: Evaluating Implicit Personalization in Dialogue Systems with Synthetic Data</i> Bo Peng, Zhiheng Wang, Heyang Gong and Chaochao Lu	17007
<i>Zero-shot Graph Reasoning via Retrieval Augmented Framework with LLMs</i> Hanqing Li, Sharika Mahadevan, Kiran Jyothi Sheena, Henry Liang and Diego Klabjan . . .	17041

<i>Privacy in Action: Towards Realistic Privacy Mitigation and Evaluation for LLM-Powered Agents</i> Shouju Wang, Fenglin Yu, Xirui Liu, Xiaoting Qin, Jue Zhang, Qingwei Lin, Dongmei Zhang and Saravan Rajmohan	17055
<i>Dissecting Logical Reasoning in LLMs: A Fine-Grained Evaluation and Supervision Study</i> Yujun Zhou, Jiayi Ye, Zipeng Ling, Yufei Han, Yue Huang, Haomin Zhuang, Zhenwen Liang, Kehan Guo, Taicheng Guo, Xiangqi Wang and Xiangliang Zhang	17075
<i>ConciseRL: Conciseness-Guided Reinforcement Learning for Efficient Reasoning Models</i> Razvan-Gabriel Dumitru, Darius Peteleaza, Vikas Yadav and Liangming Pan	17099
<i>Faster and Better LLMs via Latency-Aware Test-Time Scaling</i> Zili Wang, Tianyu Zhang, Haoli Bai, Lu Hou, Xianzhi Yu, Wulong Liu, Shiming Xiang and Lei Zhu	17124
<i>Reasoning-Augmented Conversation for Multi-Turn Jailbreak Attacks on Large Language Models</i> Zonghao Ying, Deyue Zhang, Zonglei Jing, Yisong Xiao, Quanchen Zou, Aishan Liu, Siyuan Liang, Xiangzheng Zhang, Xianglong Liu and Dacheng Tao	17138
<i>Distilling Many-Shot In-Context Learning into a Cheat Sheet</i> Ukyo Honda, Soichiro Murakami and Peinan Zhang	17158
<i>Tracing Training Footprints: A Calibration Approach for Membership Inference Attacks Against Multi-modal Large Language Models</i> Xiaofan Zheng, Huixuan Zhang and Xiaojun Wan	17179
<i>PoBiX: Detecting LLMs' Political Bias in Fact-Checking through X-phemisms</i> Charlott Jakob, David Harbecke, Patrick Parschan, Pia Wenzel Neves and Vera Schmitt ...	17192
<i>URO-Bench: Towards Comprehensive Evaluation for End-to-End Spoken Dialogue Models</i> Ruiqi Yan, Xiquan Li, Wenxi Chen, Zhikang Niu, Chen Yang, Ziyang Ma, Kai Yu and Xie Chen	17211
<i>Low-Hallucination and Efficient Coreference Resolution with LLMs</i> Yujian Gan, Yuan Liang, Jinxia Xie, Yanni Lin, Juntao Yu and Massimo Poesio	17243
<i>Your Mileage May Vary: How Empathy and Demographics Shape Human Preferences in LLM Responses</i> Yishan Wang, Amanda Cercas Curry and Flor Miriam Plaza-del-Arco	17257
<i>Diving into Mitigating Hallucinations from a Vision Perspective for Large Vision-Language Models</i> Weihang Wang, Xinhao Li, Ziyue Wang, Yan Pang, Jielei Zhang, Peiyi Li, Qiang Zhang and Longwen Gao	17271
<i>PhysicsArena: The First Multimodal Physics Reasoning Benchmark Exploring Variable, Process, and Solution Dimensions</i> Song Dai, Yibo Yan, Jiamin Su, Zihao Dongfang, Yubo Gao, Yonghua Hei, Jungang Li, Junyan Zhang, Sicheng Tao, Zhuoran Gao and Xuming Hu	17290
<i>Ko-LongRAG: A Korean Long-Context RAG Benchmark Built with a Retrieval-Free Approach</i> Yongil Kim, Heuiyeon Yeen, Hyeongu Yun and Jinsik Lee	17317
<i>Choosing a Model, Shaping a Future: Comparing LLM Perspectives on Sustainability and its Relationship with AI</i> Annika Bush, Meltem Aksoy, Markus Pauly and Greta Ontrup	17330

<i>Optimising Factual Consistency in Summarisation via Preference Learning from Multiple Imperfect Metrics</i>	
Yuxuan Ye, Raul Santos-Rodriguez and Edwin Simpson	17342
<i>Judging with Many Minds: Do More Perspectives Mean Less Prejudice? On Bias Amplification and Resistance in Multi-Agent Based LLM-as-Judge</i>	
Chiyu Ma, Enpei Zhang, Yilun Zhao, Wenjun Liu, Yaning Jia, Peijun Qing, Lin Shi, Arman Cohan, Yujun Yan and Soroush Vosoughi	17356
<i>Investigating the Impact of Conceptual Metaphors on LLM-based NLI through Shapley Interactions</i>	
Meghdut Sengupta, Maximilian Muschalik, Fabian Fumagalli, Barbara Hammer, Eyke Hüllermeier, Debanjan Ghosh and Henning Wachsmuth	17393
<i>KurTail : Kurtosis-based LLM Quantization</i>	
Mohammad Sadegh Akhondzadeh, Aleksandar Bojchevski, Evangelos Eleftheriou and Martino Dazzi	17404
<i>VIVA+: Human-Centered Situational Decision-Making</i>	
Zhe Hu, Yixiao Ren, Guanzhong Liu, Jing Li and Yu Yin	17420
<i>QuantAgents: Towards Multi-agent Financial System via Simulated Trading</i>	
Xiangyu Li, Yawen Zeng, Xiaofen Xing, Jin Xu and Xiangmin Xu	17438
<i>LLMs Reproduce Stereotypes of Sexual and Gender Minorities</i>	
Ruby Ostrow and Adam Lopez	17465
<i>Accept or Deny? Evaluating LLM Fairness and Performance in Loan Approval across Table-to-Text Serialization Approaches</i>	
Israel Abebe Azime, Deborah D. Kanubala, Tejumade Afonja, Mario Fritz, Isabel Valera, Dietrich Klakow and Philipp Slusallek	17478
<i>Transfer-Aware Data Selection for Domain Adaptation in Text Retrieval</i>	
Linzhu Yu, Huan Li, Ke Chen and Lidan Shou	17504
<i>Understanding and Improving Information Preservation in Prompt Compression for LLMs</i>	
Weronika Łajewska, Momchil Hardalov, Laura Aina, Neha Anna John, Hang Su and Lluís Marquez	17520
<i>A Benchmark for Hindi Verb-Argument Structure Alternations</i>	
Kanishka Jain and Ashwini Vaidya	17542
<i>Beyond Binary Preferences: Semi-Online Label-Free GRACE-KTO with Group-Wise Adaptive Calibration for High-Quality Long-Text Generation</i>	
Jingyang Deng, Ran Chen, Jo-Ku Cheng and Jinwen Ma	17550
<i>Representation-based Broad Hallucination Detectors Fail to Generalize Out of Distribution</i>	
Zuzanna Dubanowska, Maciej Żelaszczyk, Michał Brzozowski, Paolo Mandica and Michał P. Karpowicz	17563
<i>MAFMO: Multi-modal Adaptive Fusion with Meta-template Optimization for Vision-Language Models</i>	
Mingrui Xie, Lulu Xu and Junliang Du	17576
<i>Multimodal UNcommonsense: From Odd to Ordinary and Ordinary to Odd</i>	
Yejin Son, Saejin Kim, Dongjun Min and Youngjae Yu	17586
<i>Analyzing Gambling Addictions: A Spanish Corpus for Understanding Pathological Behavior</i>	
Manuel Couto, Marcos Fernández-Pichel, Mario Ezra Aragon and David E. Losada	17610

<i>Beyond Surface Alignment: Rebuilding LLMs Safety Mechanism via Probabilistically Ablating Refusal Direction</i>	
Yuanbo Xie, Yingjie Zhang, Tianyun Liu, Duohe Ma and Tingwen Liu	17620
<i>Distributed LLM Serving on Consumer-Grade GPUs by Reconciling Computation and Communication</i>	
Lewei Jin, Kui Zhang, Yongqi Chen, Zhuoyifan, Renjie Li, Yi Gao, Bowei Yang, Zhengong Cai and Wei Dong	17633
<i>SafeToolBench: Pioneering a Prospective Benchmark to Evaluating Tool Utilization Safety in LLMs</i>	
Hongfei Xia, Hongru Wang, Zeming Liu, Qian Yu, Yuhang Guo and Haifeng Wang	17643
<i>Sparsifying Mamba</i>	
An Wang, Ruobing Xie, Shuaipeng Li, Xingwu Sun and Zhanhui Kang	17661
<i>Beneath the Facade: Probing Safety Vulnerabilities in LLMs via Auto-Generated Jailbreak Prompts</i>	
Heehyeon Kim, Kyeongryul Lee and Joyce Jiyoungh Whang	17668
<i>ET-MIER: Entity Type-guided Key Mention Identification and Evidence Retrieval for Document-level Relation Extraction</i>	
Xin Li, Huangming Xu, Fu Zhang and Jingwei Cheng	17701
<i>Position IDs Matter: An Enhanced Position Layout for Efficient Context Compression in Large Language Models</i>	
Runsong Zhao, Xin Liu, Xinyu Liu, Pengcheng Huang, Chunyang Xiao, Tong Xiao and JingBo Zhu	17715
<i>Can Role Vectors Affect LLM Behaviour?</i>	
Daniele Poterì, Andrea Seveso and Fabio Mercorio	17735
<i>Semantic Component Analysis: Introducing Multi-Topic Distributions to Clustering-Based Topic Modeling</i>	
Florian Eichin, Carolin M. Schuster, Georg Groh and Michael A. Hedderich	17748
<i>ThinkQE: Query Expansion via an Evolving Thinking Process</i>	
Yibin Lei, Tao Shen and Andrew Yates	17772
<i>Hierarchical Reward Modeling for Fault Localization in Large Code Repositories</i>	
Jiwei Zhang, Jianxun Lian, Haiming Qin, Mingyang Zhou, KeZhong Lu, Rui Mao and Hao Liao	17782
<i>Layer Duplication in LLMs</i>	
Neo Eyal, Nachum Dershowitz and Kfir Bar	17797
<i>Semantic-Aware Action Space Compression via LLM-DRL Synergy for Efficient Task-oriented Dialogue Policy Exploration</i>	
Yangyang Zhao, Ben Niu, Yuxuan Tan, Shihan Wang and Libo Qin	17808
<i>Linear Steerability in Language Models: When It Emerges and How It Evolves</i>	
Jianshu She, Xinyue Li, Eric P. Xing, Zhengzhong Liu and Qirong Ho	17821
<i>A Comprehensive Survey on Learning from Rewards for Large Language Models: Reward Models and Learning Strategies</i>	
Xiaobao Wu	17847
<i>InFact: Informativeness Alignment for Improved LLM Factuality</i>	
Roi Cohen, Russa Biswas and Gerard de Melo	17876

<i>Large Language Model Agents in Finance: A Survey Bridging Research, Practice, and Real-World Deployment</i>	
Yifei Dong, Fengyi Wu, Kunlin Zhang, Yilong Dai, Sanjian Zhang, Wanghao Ye, Sihan Chen and Zhi-Qi Cheng	17889
<i>Problem Solved? Information Extraction Design Space for Layout-Rich Documents using LLMs</i>	
Gaye Colakoglu, Gürkan Solmaz and Jonathan Fürst	17908
<i>Generation-Augmented Retrieval: Rethinking the Role of Large Language Models in Zero-Shot Relation Extraction</i>	
Zehan Li, Fu Zhang, Tianyue Peng, He Liu and Jingwei Cheng	17928
<i>Following Occam's Razor: Dynamic Combination of Structured Knowledge for Multi-Hop Question Answering using LLMs</i>	
Wei Chen, Zhi Zheng, Lili Zhao, Huijun Hou and Tong Xu	17942
<i>Large Language Models as Reader for Bias Detection</i>	
Xuan Luo, Jing Li, Zhong Wenzhong, Geng Tu and Ruifeng Xu	17957
<i>LOHRec: Leveraging Order and Hierarchy in Generative Sequential Recommendation</i>	
Jiawen Xie, Haiyang Wu, Deyi Ji, Yuekui Yang and Shaoping Ma	17968
<i>Biology-Instructions: A Dataset and Benchmark for Multi-Omics Sequence Understanding Capability of Large Language Models</i>	
Haonan he, Yuchen Ren, Yining Tang, Ziyang Xu, Junxian Li, Minghao Yang, Di Zhang, Yuan Dong, Tao Chen, Shufei Zhang, Yuqiang Li, Nanqing Dong, Wanli Ouyang, Dongzhan Zhou and Peng Ye	17984
<i>AssistedDS: Benchmarking How External Domain Knowledge Assists LLMs in Automated Data Science</i>	
An Luo, Xun Xian, Jin Du, Fangqiao Tian, Ganghua Wang, Ming Zhong, Shengchun Zhao, Xuan Bi, Zirui Liu, Jiawei Zhou, Jayanth Srinivasa, Ashish Kundu, Charles Fleming, Mingyi Hong and Jie Ding	18017
<i>Are you sure? Measuring models bias in content moderation through uncertainty</i>	
Alessandra Urbinati, Mirko Lai, Simona Frenda and Marco Stranisci	18061
<i>FOSSIL: Harnessing Feedback on Suboptimal Samples for Data-Efficient Generalisation with Imitation Learning for Embodied Vision-and-Language Tasks</i>	
Sabrina McCallum, Amit Parekh and Alessandro Suglia	18077
<i>Assess and Prompt: A Generative RL Framework for Improving Engagement in Online Mental Health Communities</i>	
Bhagesh Gaur, Karan Gupta, Aseem Srivastava, Manish Gupta and Md Shad Akhtar	18102
<i>Logic: Long-form Outline Generation via Imitative and Critical Self-refinement</i>	
Hengwei Liu, Yongliang Shen, Zhe Zheng, Haoyuan Ma, Xingyu Wu, Yin Zhang and Weiming Lu	18119
<i>No Free Lunch: Retrieval-Augmented Generation Undermines Fairness in LLMs, Even for Vigilant Users</i>	
Mengxuan Hu, Hongyi Wu, Ronghang Zhu, Zihan Guan, Dongliang Guo, Daiqing Qi and Sheng Li	18145
<i>LegoSLM: Connecting LLM with Speech Encoder using CTC Posteriors</i>	
Rao Ma, Tongzhou Chen, Kartik Audhkhasi and Bhuvana Ramabhadran	18171

<i>Skeleton-Guided-Translation: A Benchmarking Framework for Code Repository Translation with Fine-Grained Quality Evaluation</i>	
Xing Zhang, Jiaheng Wen, Fangkai Yang, Yu Kang, Pu Zhao, Junhao Wang, Maoquan Wang, Yufan Huang, Shengyu Fu, Elsie Nallipogu, Qingwei Lin, Yingnong Dang, Saravan Rajmohan and Dongmei Zhang	18187
<i>Parallel Communities Across the Surface Web and the Dark Web</i>	
Wenchao Dong, Megha Sundriyal, Seongchan Park, Jaehong Kim, Meeyoung Cha, Tanmoy Chakraborty and Wonjae Lee	18199
<i>Lemma Dilemma: On Lemma Generation Without Domain- or Language-Specific Training Data</i>	
Olia Toporkov, Alan Akbik and Rodrigo Agerri	18219
<i>LlmFixer: Fix the Helpfulness of Defensive Large Language Models</i>	
Zelong Yu, Xiaoming Zhang, Litian Zhang, Yu Yuan and Chaozhuo Li	18233
<i>Universal Acoustic Adversarial Attacks for Flexible Control of Speech-LLMs</i>	
Rao Ma, Mengjie Qian, Vyas Raina, Mark Gales and Kate Knill	18248
<i>Probing Semantic Routing in Large Mixture-of-Expert Models</i>	
Matthew Lyle Olson, Neale Ratzlaff, Musashi Hinck, Man Luo, Sungduk Yu, Chendi Xue and Vasudev Lal	18263
<i>CMT-Eval: A Novel Chinese Multi-turn Dialogue Evaluation Dataset Addressing Real-world Conversational Challenges</i>	
Siyu Tian, Kaijie Mo, Yupei Wang and Renfen Hu	18279
<i>LastingBench: Defend Benchmarks Against Knowledge Leakage</i>	
Yixiong Fang, Tianran Sun, Yuling Shi, Min Wang and Xiaodong Gu	18304
<i>Learning API Functionality from In-Context Demonstrations for Tool-based Agents</i>	
Bhrij Patel, Ashish Jagmohan and Aditya Vempaty	18318
<i>Predicting Language Models' Success at Zero-Shot Probabilistic Prediction</i>	
Kevin Ren, Santiago Cortes-Gomez, Carlos Miguel Patiño, Ananya Joshi, Ruiqi Lyu, Jingjing Tang, Alistair Turcan, Khurram Yamin, Steven Wu and Bryan Wilder	18337
<i>GAMIC: Graph-Aligned Molecular In-context Learning for Molecule Analysis via LLMs</i>	
Ali Al Lawati, Jason S Lucas, Zhiwei Zhang, Prasenjit Mitra and Suhang Wang	18364
<i>Rethinking Sign Language Translation: The Impact of Signer Dependence on Model Evaluation</i>	
Keren Artiaga, Sabyasachi Kamila, Haithem Afli, Conor Lynch and Mohammed Hasanuzzaman	18379
<i>Can Large Language Models Identify Implicit Suicidal Ideation? An Empirical Evaluation</i>	
Tong Li, Shu Yang, Junchao Wu, Jiayao Wei, Lijie Hu, Mengdi Li, Derek F. Wong, Joshua R. Oltmanns and Di Wang	18392
<i>Adaptive Platt Scaling with Causal Interpretations for Self-Reflective Language Model Uncertainty Estimates</i>	
Anthony Sicilia and Malihe Alikhani	18414
<i>Treble Counterfactual VLMs: A Causal Approach to Hallucination</i>	
Li Li, Jiashu Qu, Linxin Song, Yuxiao Zhou, Yuehan Qin, Tiankai Yang and Yue Zhao ...	18423
<i>Video-Skill-CoT: Skill-based Chain-of-Thoughts for Domain-Adaptive Video Reasoning</i>	
Daeun Lee, Jaehong Yoon, Jaemin Cho and Mohit Bansal	18435

<i>Glitter: A Multi-Sentence, Multi-Reference Benchmark for Gender-Fair German Machine Translation</i> A Pranav, Janiça Hackenbuchner, Giuseppe Attanasio, Manuel Lardelli and Anne Lauscher	18450
<i>From n-gram to Attention: How Model Architectures Learn and Propagate Bias in Language Modeling</i> Mohsinul Kabir, Tasfia Tahsin and Sophia Ananiadou	18478
<i>SENTRA: Selected-Next-Token Transformer for LLM Text Detection</i> Mitchell Plyler, Yilun Zhang, Alexander Tuzhilin, Saoud Khalifah and Sen Tian	18499
<i>Automate Strategy Finding with LLM in Quant Investment</i> Zhizhuo Kou, Holam Yu, Junyu Luo, Jingshu Peng, Xujia Li, Chengzhong Liu, Juntao Dai, Lei Chen, Sirui Han and Yike Guo	18517
<i>Does Reasoning Introduce Bias? A Study of Social Bias Evaluation and Mitigation in LLM Reasoning</i> Xuyang Wu, Jinming Nian, Ting-Ruen Wei, Zhiqiang Tao, Hsin-Tai Wu and Yi Fang	18534
<i>MT-RewardTree: A Comprehensive Framework for Advancing LLM-Based Machine Translation via Reward Modeling</i> Zhaopeng Feng, Jiahan Ren, Jiayuan Su, Jiamei Zheng, Hongwei Wang and Zuozhu Liu	18556
<i>Bias after Prompting: Persistent Discrimination in Large Language Models</i> Nivedha Sivakumar, Natalie Mackraz, Samira Khorshidi, Krishna Patel, Barry-John Theobald, Luca Zappella and Nicholas Apostoloff	18568
<i>CARVQ: Corrective Adaptor with Group Residual Vector Quantization for LLM Embedding Compression</i> Dayin Gou, Sanghyun Byun, Nilesh Malpeddi, Gabrielle De Micheli, Prathamesh Vaste, Jacob Song and Woo Seong Chung	18594
<i>Consistent Discourse-level Temporal Relation Extraction Using Large Language Models</i> Yi Fan and Michael Strube	18605
<i>MMPlanner: Zero-Shot Multimodal Procedural Planning with Chain-of-Thought Object State Reasoning</i> Afrina Tabassum, Bin Guo, Xiyao Ma, Hoda Eldardiry and Ismini Lourentzou	18623
<i>Internal states before wait modulate reasoning patterns</i> Dmitrii Troitskii, Koyena Pal, Chris Wendler and Callum Stuart McDougall	18640
<i>Sparsity May Be All You Need: Sparse Random Parameter Adaptation</i> Jesus Rios, Pierre Dognin, Ronny Luss and Karthikeyan Natesan Ramamurthy	18650
<i>Learning to Align: Addressing Character Frequency Distribution Shifts in Handwritten Text Recognition</i> Panagiotis Kaliosis and John Pavlopoulos	18667
<i>MT-R1-Zero: Advancing LLM-based Machine Translation via R1-Zero-like Reinforcement Learning</i> Zhaopeng Feng, Shaosheng Cao, Jiahan Ren, Jiayuan Su, Ruizhe Chen, Yan Zhang, Jian Wu and Zuozhu Liu	18685
<i>Discrete Minds in a Continuous World: Do Language Models Know Time Passes?</i> Minghan Wang, Ye Bai, Thuy-Trang Vu, Ehsan Shareghi and Gholamreza Haffari	18703
<i>DLTKG: Denoising Logic-based Temporal Knowledge Graph Reasoning</i> Xiaoke Wang, Fu Zhang, Jingwei Cheng, Yiwen Chi, Jiashun Peng and Yingsong Ning	18730

<i>EMO-RL: Emotion-Rule-Based Reinforcement Learning Enhanced Audio-Language Model for Generalized Speech Emotion Recognition</i>	
Pengcheng Li, Botao Zhao, Zuheng Kang, Junqing Peng, Xiaoyang Qu, Yayun He and Jianzong Wang	18744
<i>MANTA: A Scalable Pipeline for Transmuting Massive Web Corpora into Instruction Datasets</i>	
Heuiyeen Yeen, Seokhee Hong, Hyeongu Yun and Jinsik Lee	18755
<i>Fast Quiet-STaR: Thinking Without Thought Tokens</i>	
Wei Huang, Yizhe Xiong, Xin Ye, Zhijie Deng, Hui Chen, Zijia Lin and Guiguang Ding ..	18771
<i>Lock on Target! Precision Unlearning via Directional Control</i>	
Yuntao Wen, Ruixiang Feng, Feng Guo, Yifan Wang, Ran Le, Yang Song, Shen Gao and Shuo Shang	18782
<i>UniRAG: A Unified RAG Framework for Knowledge-Intensive Queries with Decomposition, Break-Down Reasoning, and Iterative Rewriting</i>	
Gun Il Kim, Jong Wook Kim and Beakcheol Jang	18795
<i>One Shot Dominance: Knowledge Poisoning Attack on Retrieval-Augmented Generation Systems</i>	
Zhiyuan Chang, Mingyang Li, Xiaojun Jia, Junjie Wang, Yuekai Huang, Ziyong Jiang, Yang Liu and Qing Wang	18811
<i>From Generic Empathy to Personalized Emotional Support: A Self-Evolution Framework for User Preference Alignment</i>	
Jing Ye, Lu Xiang, Yaping Zhang and Chengqing Zong	18826
<i>MaskCD: Mitigating LLM Hallucinations by Image Head Masked Contrastive Decoding</i>	
Jingyuan Deng and Yujiu Yang	18854
<i>ClusterUCB: Efficient Gradient-Based Data Selection for Targeted Fine-Tuning of LLMs</i>	
Zige Wang, Qi Zhu, Fei Mi, Minghui Xu, Ruochun Jin and Wenjing Yang	18867
<i>TrapDoc: Deceiving LLM Users by Injecting Imperceptible Phantom Tokens into Documents</i>	
Hyundong Jin, Sicheol Sung, Shinwoo Park, SeungYeop Baik and Yo-Sub Han	18881
<i>AraReasoner: Evaluating Reasoning-Based LLMs for Arabic NLP</i>	
Ahmed Abul Hasanaath, Aisha Alansari, Ahmed Ashraf, Salmane Chafik, Hamzah Luqman and Saad Ezzini	18898
<i>Tales of Morality: Comparing Human- and LLM-Generated Moral Stories from Visual Cues</i>	
Rezvaneh Rezapour, Sullam Jeoung, Zhiwen You and Jana Diesner	18915
<i>AirRAG: Autonomous Strategic Planning and Reasoning Steer Retrieval Augmented Generation</i>	
Wenfeng Feng, Chuzhan Hao, Yuewei Zhang, Guochao Jiang and Jingyi Song	18934
<i>Evaluating NL2SQL via SQL2NL</i>	
Mohammadtaher Safarzadeh, Afshin Oroojlooy and Dan Roth	18954
<i>DB-Explore: Automated Database Exploration and Instruction Synthesis for Text-to-SQL</i>	
Haoyuan Ma, Yongliang Shen, Hengwei Liu, Wenqi Zhang, Haolei Xu, Qiuying Peng, Jun Wang and Weiming Lu	18969
<i>Do BERT-Like Bidirectional Models Still Perform Better on Text Classification in the Era of LLMs?</i>	
Junyan Zhang, Yiming Huang, Shuliang Liu, Yubo Gao and Xuming Hu	18980

<i>Divide, Optimize, Merge: Scalable Fine-Grained Generative Optimization for LLM Agents</i>	
Jiale Liu, Yifan Zeng, Shaokun Zhang, Chi Zhang, Malte Højmark-Bertelsen, Marie Normann Gadeberg, Huazheng Wang and Qingyun Wu	18990
<i>Evaluating Evaluation Metrics – The Mirage of Hallucination Detection</i>	
Atharva Kulkarni, Yuan Zhang, Joel Ruben Antony Moniz, Xiou Ge, Bo-Hsiang Tseng, Dhivya Piraviperumal, Swabha Swayamdipta and Hong Yu	19013
<i>The Progress Illusion: Revisiting meta-evaluation standards of LLM evaluators</i>	
Tianruo Rose Xu, Vedant Gaur, Liu Leqi and Tanya Goyal	19033
<i>MidPO: Dual Preference Optimization for Safety and Helpfulness in Large Language Models via a Mixture of Experts Framework</i>	
Yupeng Qi, Ziyu Lyu, Min Yang, Yanlin Wang, Lu Bai and Lixin Cui	19044
<i>From KMMLU-Redux to Pro: A Professional Korean Benchmark Suite for LLM Evaluation</i>	
Seokhee Hong, Sunkyoung Kim, Guijin Son, Soyeon Kim, Yeonjung Hong and Jinsik Lee	19067
<i>RealBench: A Chinese Multi-image Understanding Benchmark Close to Real-world Scenarios</i>	
Fei Zhao, Chengqiang Lu, Yufan Shen, Qimeng Wang, Yicheng Qian, Haoxin Zhang, Yan Gao, Yiwu, Yao Hu, Zhen Wu, Shangyu Xing and Xinyu Dai	19097
<i>The More, The Better? A Critical Study of Multimodal Context in Radiology Report Summarization</i>	
Mong Yuan Sim, Wei Emma Zhang, Xiang Dai, Biaoyan Fang, Sarbin Ranjitkar, Arjun Burlakoti, Jamie Taylor and Haojie Zhuang	19116
<i>Localizing Malicious Outputs from CodeLLM</i>	
Mayukh Borana, Junyi Liang, Sai Sathiesh Rajan and Sudipta Chattopadhyay	19132
<i>Knowing More, Acting Better: Hierarchical Representation for Embodied Decision-Making</i>	
Chunhui Zhang, Zhongyu Ouyang, Xingjian Diao, Zheyuan Liu and Soroush Vosoughi . . .	19144
<i>Culture is Everywhere: A Call for Intentionally Cultural Evaluation</i>	
Juhyun Oh, Inha Cha, Michael Saxon, Hyunseung Lim, Shaily Bhatt and Alice Oh	19156
<i>Fairness in Automatic Speech Recognition Isn't a One-Size-Fits-All</i>	
Hend ElGhazaly, Bahman Mirheidari, Heidi Christensen and Nafise Sadat Moosavi	19169
<i>Uncovering Factor-Level Preference to Improve Human-Model Alignment</i>	
Juhyun Oh, Eunsu Kim, Jiseon Kim, Wenda Xu, Inha Cha, William Yang Wang and Alice Oh	19179
<i>Adaptive Preference Optimization with Uncertainty-aware Utility Anchor</i>	
Xiaobo Wang, Zixia Jia, Jiaqi Li, Qi Liu and Zilong Zheng	19204
<i>GRAD: Generative Retrieval-Aligned Demonstration Sampler for Efficient Few-Shot Reasoning</i>	
Oussama Gabouj, Kamel Charaf, Ivan Zakazov, Nicolas Baldwin and Robert West	19226
<i>IoTmigrator: LLM-driven Embedded IoT Code Migration across Different OSes for Cloud-device Integration</i>	
YQ, Kaijie Gong, Yi Gao, Hao Wang and Wei Dong	19245
<i>ClueAnchor: Clue-Anchored Knowledge Reasoning Exploration and Optimization for Retrieval-Augmented Generation</i>	
Hao Chen, Yukun Yan, Sen Mei, Wanxiang Che, Zhenghao Liu, Qi Shi, Xinze Li, Yuchun Fan, Pengcheng Huang, Qiushi Xiong, Zhiyuan Liu and Maosong Sun	19258

<i>BAGELS: Benchmarking the Automated Generation and Extraction of Limitations from Scholarly Text</i> Ibrahim Al Azher, Miftahul Jannat Mokarrama, Zhishuai Guo, Sagnik Ray Choudhury and Hamed Alhoori	19279
<i>Dense Retrievers Can Fail on Simple Queries: Revealing The Granularity Dilemma of Embeddings</i> Liyan Xu, Zhenlin Su, Mo Yu, Jiangnan Li, Fandong Meng and Jie Zhou	19295
<i>Over-Generation and Compaction: A Prompting Strategy for Procedural Text Adaptation with Large Language Models</i> HyeongSik Kim, XU Yanheng, Chaoqun Dong and Fei Du	19306
<i>TransBERT: A Framework for Synthetic Translation in Domain-Specific Language Modeling</i> Julien Knafou, Luc Mottin, Anaïs Mottaz, Alexandre Flament and Patrick Ruch	19338
<i>Beyond Fixed-Length Calibration for Post-Training Compression of LLMs</i> Jaehoon Oh and Dokwan Oh	19355
<i>Attributes as Textual Genes: Leveraging LLMs as Genetic Algorithm Simulators for Conditional Synthetic Data Generation</i> Guangzeng Han, Weisi Liu and Xiaolei Huang	19367
<i>ReCoVeR the Target Language: Language Steering without Sacrificing Task Performance</i> Hannah Sterz, Fabian David Schmidt, Goran Glavaš and Ivan Vulić	19390
<i>LC-Eval: A Bilingual Multi-Task Evaluation Benchmark for Long-Context Understanding</i> Sheikh Jubair, Arwa Omayrah, Amal Alshammari, Alhanoof Althnian, Abdulhamed Alothaimen, Norah A. Alzahrani, Shahad D. Alzaidi, Nora Al-Twairish and Abdulmohsen Al-Thubaity	19406
<i>OVFact: Measuring and Improving Open-Vocabulary Factuality for Long Caption Models</i> Monika Wysoczańska, Shyamal Buch, Anurag Arnab and Cordelia Schmid	19440
<i>GRPO-Guided Modality Selection Enhanced LoRA-Tuned LLMs for Multimodal Emotion Recognition</i> Yang Chen, Shuwan Yang, Yan Xiang, Ran Song, Yuxin Huang and Zhengtao Yu	19458
<i>Defending against Indirect Prompt Injection by Instruction Detection</i> Tongyu Wen, Chenglong Wang, Xiyuan Yang, Haoyu Tang, Yueqi Xie, Lingjuan Lyu, Zhicheng Dou and Fangzhao Wu	19472
<i>MUG-Eval: A Proxy Evaluation Framework for Multilingual Generation Capabilities in Any Language</i> Seyoung Song, Seogyong Jeong, Eunsu Kim, Jiho Jin, Dongkwan Kim, Jay Shin and Alice Oh	19488
<i>CAC-CoT: Connector-Aware Compact Chain-of-Thought for Efficient Reasoning Data Synthesis Across Dual-System Cognitive Tasks</i> Sunguk Choi, Yonghoon Kwon and Heondeuk Lee	19515
<i>On the Versatility of Sparse Autoencoders for In-Context Learning</i> Ikhyun Cho, Gaeul Kwon and Julia Hockenmaier	19531
<i>More Documents, Same Length: Isolating the Challenge of Multiple Documents in RAG</i> Shahar Levy, Nir Mazor, Lihi Shalmon, Michael Hassid and Gabriel Stanovsky	19539
<i>CLEAR: A Comprehensive Linguistic Evaluation of Argument Rewriting by Large Language Models</i> Thomas Huber and Christina Niklaus	19548
<i>ALRPHFS: Adversarially Learned Risk Patterns with Hierarchical Fast & Slow Reasoning for Robust Agent Defense</i> Shiyu Xiang, Tong Zhang and Ronghao Chen	19569

<i>Stop Playing the Guessing Game! Evaluating Conversational Recommender Systems via Target-free User Simulation</i>	
SungHwan Kim, Kwangwook Seo, Tongyoung Kim, Jinyoung Yeo and Dongha Lee	19588
<i>Out-of-Context Reasoning in Large Language Models</i>	
Jonathan Shaki, Emanuele La Malfa, Michael J. Wooldridge and Sarit Kraus	19606
<i>CodeComplex: Dataset for Worst-Case Time Complexity Prediction</i>	
SeungYeop Baik, Joonghyuk Hahn, Jungin Kim, Aditi , Mingi Jeon, Yo-Sub Han and Sang-Ki Ko	19616
<i>Weak2Wise: An Automated, Lightweight Framework for Weak-LLM-Friendly Reasoning Synthesis</i>	
Jianing Lin, Yuanfang Guo, Shunning Liu, Zeming Liu and Yunhong Wang	19639
<i>From Tower to Spire: Adding the Speech Modality to a Translation-Specialist LLM</i>	
Kshitij Ambilduke, Ben Peters, Sonal Sannigrahi, Anil Keshwani, Tsz Kin Lam, Bruno Martins, Andre Martins and Marcely Zanon Boito	19658
<i>LLM Agents at the Roundtable: A Multi-Perspective and Dialectical Reasoning Framework for Essay Scoring</i>	
Jinhee Jang, Ayoung Moon, Minkyung Jung, YoungBin Kim and Seung Jin Lee	19674
<i>DeepNote: Note-Centric Deep Retrieval-Augmented Generation</i>	
Ruobing Wang, Qingfei Zhao, Yukun Yan, Daren Zha, Yuxuan Chen, Shi Yu, Zhenghao Liu, Yixuan Wang, Shuo Wang, Xu Han, Zhiyuan Liu and Maosong Sun	19688
<i>NormAL LoRA: What is the perfect size?</i>	
Aastik, Topu Sai Meghana, Chinmay Prakash Kulkarni and Pragya Paramita Sahu	19716
<i>Inclusive Leadership in the Age of AI: A Dataset and Comparative Study of LLMs vs. Real-Life Leaders in Workplace Action Planning</i>	
Vindhya Singh, Sabine Schulte im Walde and Ksenia Keplinger	19732
<i>Token Preference Optimization with Self-Calibrated Visual-Anchored Rewards for Hallucination Mitigation</i>	
Jihao Gu, Yingyao Wang, Meng Cao, Pi Bu, Jun Song, Bo Zheng, Yancheng He and Shilong Li	19754
<i>EZ-VC: Easy Zero-shot Any-to-Any Voice Conversion</i>	
Advait Joglekar, Divyanshu Singh, Rooshil Rohit Bhatia and Srinivasan Umesh	19768
<i>Length Representations in Large Language Models</i>	
Sangjun Moon, Dasom Choi, Jingun Kwon, Hidetaka Kamigaito and Manabu Okumura . .	19775
<i>MultiLingPoT: Boosting Mathematical Reasoning in LLMs through Multilingual Program Integration</i>	
Nianqi Li, Zujie Liang, Siyu Yuan, Jiaqing Liang, Feng Wei and Yanghua Xiao	19794
<i>Simulating Identity, Propagating Bias: Abstraction and Stereotypes in LLM-Generated Text</i>	
Pia Sommerauer, Giulia Rambelli and Tommaso Caselli	19812
<i>Do LVLMs Know What They Know? A Systematic Study of Knowledge Boundary Perception in LVLMs</i>	
Zhikai Ding, Shiyu Ni and Keping Bi	19832
<i>Benchmarking Large Language Models for Cryptanalysis and Side-Channel Vulnerabilities</i>	
Utsav Maskey, Chencheng Zhu and Usman Naseem	19849
<i>MTabVQA: Evaluating Multi-Tabular Reasoning of Language Models in Visual Space</i>	
Anshul Singh, Chris Biemann and Jan Strich	19866

<i>TurnBench-MS: A Benchmark for Evaluating Multi-Turn, Multi-Step Reasoning in Large Language Models</i>	Yiran Zhang, Mo Wang, Xiaoyang Li, Kaixuan Ren, Chencheng Zhu and Usman Naseem .	19892
<i>Assessing LLM Reasoning Steps via Principal Knowledge Grounding</i>	Hyeon Hwang, Yewon Cho, Chanwoong Yoon, Yein Park, Minju Song, Kyungjae Lee, Gangwoo Kim and Jaewoo Kang	19925
<i>Stratified Selective Sampling for Instruction Tuning with Dedicated Scoring Strategy</i>	Paramita Mirza, Lucas Weber and Fabian K�uch	19949
<i>CoTD-PO: Chain-of-Thought Distillation with Preference Optimization</i>	Lujie Niu, Haochen Sun, Fangkun Zhao, Sheng Chen, Zimeng Bai, Jiawei Zhang, Caixia Yuan and Xiaojie Wang	19975
<i>Intelligent Document Parsing: Towards End-to-end Document Parsing via Decoupled Content Parsing and Layout Grounding</i>	Hangdi Xing, Feiyu Gao, Qi Zheng, Zhaoqing Zhu, Zirui Shao and Ming Yan	19987
<i>Feel the Difference? A Comparative Analysis of Emotional Arcs in Real and LLM-Generated CBT Sessions</i>	Xiaoyi Wang, Jiwei Zhang, Guangtao Zhang and Honglei Guo	19999
<i>Beyond Single-User Dialogue: Assessing Multi-User Dialogue State Tracking Capabilities of Large Language Models</i>	Sangmin Song, Juhwan Choi, JungMin Yun and YoungBin Kim	20018
<i>All-in-one: Understanding and Generation in Multimodal Reasoning with the MAIA Benchmark</i>	Davide Testa, Giovanni Bonetta, Raffaella Bernardi, Alessandro Bondielli, Alessandro Lenci, Alessio Miaschi, Lucia Passaro and Bernardo Magnini	20030
<i>Triangulating LLM Progress through Benchmarks, Games, and Cognitive Tests</i>	Filippo Moment�, Alessandro Suglia, Mario Giulianelli, Ambra Ferrari, Alexander Koller, Oliver Lemon, David Schlangen, Raquel Fern�andez and Raffaella Bernardi	20051
<i>Entity Profile Generation and Reasoning with LLMs for Entity Alignment</i>	Rumana Ferdous Munne, Md Mostafizur Rahman and Yuji Matsumoto	20073
<i>Re-FRAME the Meeting Summarization SCOPE: Fact-Based Summarization and Personalization via Questions</i>	Frederic Kirstein, Sonu Kumar, Terry Ruas and Bela Gipp	20087
<i>Attack as Defense: Safeguarding Large Vision-Language Models from Jailbreaking by Adversarial Attacks</i>	Chongxin Li, Hanzhang Wang and Yuchun Fang	20138
<i>Emphasising Structured Information: Integrating Abstract Meaning Representation into LLMs for Enhanced Open-Domain Dialogue Evaluation</i>	Bohao Yang, Kun Zhao, Dong Liu, Chen Tang, Liang Zhan and Chenghua Lin	20153
<i>Differentiated Vision: Unveiling Entity-Specific Visual Modality Requirements for Multimodal Knowledge Graph</i>	Minghang Liu, Yinghan Shen, Zihe Huang, Yuanzhuo Wang, Xuhui Jiang and Huawei Shen	20170
<i>Post Persona Alignment for Multi-Session Dialogue Generation</i>	Yi-Pei Chen, Noriki Nishida, Hideki Nakayama and Yuji Matsumoto	20184

<i>MASSIVE-Agents: A Benchmark for Multilingual Function-Calling in 52 Languages</i> Mayank Kulkarni, Vittorio Mazzia, Judith Gaspers, Chris Hench and Jack FitzGerald	20193
<i>Crafting Customisable Characters with LLMs: A Persona-Driven Role-Playing Agent Framework</i> Bohao Yang, Dong Liu, Chenghao Xiao, Kun Zhao, Chen Tang, Chao Li, Lin Yuan, Yang Guang and Chenghua Lin	20216
<i>Can LLMs Express Personality Across Cultures? Introducing Cultural Personas for Evaluating Trait Alignment</i> Priyanka Dey, Aayush Bothra, Yugal Khanter, Jieyu Zhao and Emilio Ferrara	20241
<i>Exploring the Hidden Reasoning Process of Large Language Models by Misleading Them</i> Guanyu Chen, Peiyang Wang, Yizhou Jiang, Yuqian Liu, Chujie Zhao, Ying Fang, Tianren Zhang and Feng Chen	20263
<i>When Models Reason in Your Language: Controlling Thinking Language Comes at the Cost of Accuracy</i> Jirui Qi, Shan Chen, Zidi Xiong, Raquel Fernández, Danielle Bitterman and Arianna Bisazza 20279	
<i>The Role of Model Confidence on Bias Effects in Measured Uncertainties for Vision-Language Models</i> Xinyi Liu, Weiguang Wang and Hangfeng He	20297
<i>GAttention: Gated Attention for the Detection of Abusive Language</i> Horacio Jarquín Vásquez, Hugo Jair Escalante, Manuel Montes and Mario Ezra Aragon . .	20314
<i>Towards Low-Resource Alignment to Diverse Perspectives with Sparse Feedback</i> Chu Fei Luo, Samuel Dahan and Xiaodan Zhu	20330
<i>ProtoXTM: Cross-Lingual Topic Modeling with Document-Level Prototype-based Contrastive Learning</i> Seung-Won Seo and Soon-Sun Kwon	20340
<i>One More Question is Enough, Expert Question Decomposition (EQD) Model for Domain Quantitative Reasoning</i> Mengyu Wang, Sotirios Sabanis, Miguel de Carvalho, Shay B Cohen and Tiejun Ma	20355
<i>When Punctuation Matters: A Large-Scale Comparison of Prompt Robustness Methods for LLMs</i> Mikhail Seleznyov, Mikhail Chaichuk, Gleb Ershov, Alexander Panchenko, Elena Tutubalina and Oleg Somov	20370
<i>RAR²: Retrieval-Augmented Medical Reasoning via Thought-Driven Retrieval</i> Kaishuai Xu, Wenjun Hou, Yi Cheng and Wenjie Li	20386
<i>The Security Threat of Compressed Projectors in Large Vision-Language Models</i> Yudong Zhang, Ruobing Xie, Xingwu Sun, Jiansheng Chen, Zhanhui Kang, Di Wang and Yu Wang	20397
<i>NarratEX Dataset: Explaining the Dominant Narratives in News Texts</i> Nuno Guimarães, Purificação Silvano, Ricardo Campos, Alipio Jorge, Ana Filipa Pacheco, Di- mitar Iliyanov Dimitrov, Nikolaos Nikolaidis, Roman Yangarber, Elisa Sartori, Nicolas Stefanovitch, Preslav Nakov, Jakub Piskorski and Giovanni Da San Martino	20408
<i>Radical Allomorphy: Phonological Surface Forms without Phonology</i> Salam Khalifa, Nizar Habash and Owen Rambow	20435
<i>Model Calibration for Emotion Detection</i> Mihaela Petre-Vlad, Cornelia Caragea and Florentina Hristea	20442

<i>From Benchmark to Better Embeddings: Leveraging Synonym Substitution to Enhance Multimodal Models in Ukrainian</i>	
Volodymyr Mudryi and Yurii Laba	20458
<i>Context Copying Modulation: The Role of Entropy Neurons in Managing Parametric and Contextual Knowledge Conflicts</i>	
Zineddine Tighidet, Andrea Mogini, Hedi Ben younes, Jiali Mei, Patrick Gallinari and Benjamin Piwowarski	20469
<i>A Generalizable Rhetorical Strategy Annotation Model Using LLM-based Debate Simulation and Labelling</i>	
Shiyu Ji, Farnoosh Hashemi, Joice Chen, Juanwen Pan, Weicheng Ma, Hefan Zhang, Sophia Pan, Ming Cheng, Shubham Mohole, Saeed Hassanpour, Soroush Vosoughi and Michael Macy	20482
<i>SecDecoding: Steerable Decoding for Safer LLM Generation</i>	
Jiayou Wang, Rundong Liu, Yue Hu, Huijia Wu and Zhaofeng He	20504
<i>GENUINE: Graph Enhanced Multi-level Uncertainty Estimation for Large Language Models</i>	
Tuo Wang, Adithya Kulkarni, Tyler Cody, Peter A. Beling, Yujun Yan and Dawei Zhou ...	20522
<i>ReviewEval: An Evaluation Framework for AI-Generated Reviews</i>	
Madhav Krishan Garg, Tejash Prasad, Tanmay Singhal, Chhavi Kirtani, Murari Mandal and Dhruv Kumar	20542
<i>Overcoming Black-box Attack Inefficiency with Hybrid and Dynamic Select Algorithms</i>	
Abhinay Shankar Belde, Rohit Ramkumar and Jonathan Rusert	20565
<i>GmSLM : Generative Marmoset Spoken Language Modeling</i>	
Talia Sternberg, Michael London, David Omer and Yossi Adi	20599
<i>QA-LIGN: Aligning LLMs through Constitutionally Decomposed QA</i>	
Jacob Dineen, Aswin Rrv, Qin Liu, Zhikun Xu, Xiao Ye, Ming Shen, Zhaonan Li, Shijie Lu, Chitta Baral, Muhao Chen and Ben Zhou	20619
<i>Characterizing Positional Bias in Large Language Models: A Multi-Model Evaluation of Prompt Order Effects</i>	
Patrick Schilcher, Dominik Karasin, Michael Schöpf, Haisam Saleh, Antonela Tommasel and Markus Schedl	20643
<i>You Only Use Reactive Attention Slice When Retrieving From Long Context</i>	
Yun Joon Soh, Hanxian Huang, Yuandong Tian and Jishen Zhao	20665
<i>Fine-Tuned Thoughts: Leveraging Chain-of-Thought Reasoning for Industrial Asset Health Monitoring</i>	
Shuxin Lin, Dhaval C Patel and Christodoulos Constantinides	20687
<i>CoViPAL: Layer-wise Contextualized Visual Token Pruning for Large Vision-Language Models</i>	
Zicong Tang, Ziyang Ma, Suqing Wang, Zuchao Li, Lefei Zhang, Hai Zhao, Yun Li and Qianren Wang	20701
<i>Large Language Models with Temporal Reasoning for Longitudinal Clinical Summarization and Prediction</i>	
Maya Kruse, Shiyue Hu, Nicholas Derby, Yifu Wu, Samantha Stonbraker, Bingsheng Yao, Dakuo Wang, Elizabeth M. Goldberg and Yanjun Gao	20715
<i>TransAlign: Machine Translation Encoders are Strong Word Aligners, Too</i>	
Benedikt Ebing, Christian Goldschmied and Goran Glavaš	20736

<i>Pruning Weights but Not Truth: Safeguarding Truthfulness While Pruning LLMs</i>	
Yao Fu, Runchao Li, Xianxuan Long, Haotian Yu, Xiaotian Han, Yu Yin and Pan Li	20750
<i>Augment before You Try: Knowledge-Enhanced Table Question Answering via Table Expansion</i>	
Yujian Liu, Jiabao Ji, Tong Yu, Ryan A. Rossi, Sungchul Kim, Handong Zhao, Ritwik Sinha, Yang Zhang and Shiyu Chang	20769
<i>Evaluating Large Language Models for Belief Inference: Mapping Belief Networks at Scale</i>	
Trisevgeni Papakonstantinou, Antonina Zhiteneva, Ana Yutong Ma, Derek Powell and Zachary Horne	20787
<i>Distinguishing fair from unfair compositional generalization tasks</i>	
Ahmad Jabbar, Cleo Condoravdi and Christopher Potts	20796
<i>SA-CLIP: Language Guided Image Spatial and Action Feature Learning</i>	
Guanlin Li, Wenhao Shao, Praboda Rajapaksha and Noel Crespi	20808
<i>Inefficiencies of Meta Agents for Agent Design</i>	
Batu El, Mert Yuksekgonul and James Zou	20815
<i>SCoder: Progressive Self-Distillation for Bootstrapping Small-Scale Data Synthesizers to Empower Code LLMs</i>	
Xinyu Zhang, Changzhi Zhou, Linmei Hu, Luhao Zhang, Xiancai Chen, Haomin Fu, Yang Yang and Mengdi Zhang	20825
<i>Linguistically-Controlled Paraphrase Generation</i>	
Mohamed Elgaar and Hadi Amiri	20842
<i>LAWCAT: Efficient Distillation from Quadratic to Linear Attention with Convolution across Tokens for Long Context Modeling</i>	
Zeyu Liu, Souvik Kundu, Lianghao Jiang, Anni Li, Srikanth Ronanki, Sravan Babu Bodapati, Gourav Datta and Peter Anthony Beerel	20865
<i>Analyzing Dialectical Biases in LLMs for Knowledge and Reasoning Benchmarks</i>	
Eileen Pan, Anna Seo Gyeong Choi, Maartje Ter Hoeve, Skyler Seto and Allison Koenecke	20882
<i>TreeBoN: Enhancing Inference-Time Alignment with Speculative Tree-Search and Best-of-N Sampling</i>	
Jiahao Qiu, Yifu Lu, Yifan Zeng, Jiacheng Guo, Jiayi Geng, Chenhao Zhu, Xinzhe Juan, Ling Yang, Huazheng Wang, Kaixuan Huang, Yue Wu and Mengdi Wang	20894
<i>CulturalFrames: Assessing Cultural Expectation Alignment in Text-to-Image Models and Evaluation Metrics</i>	
Shravan Nayak, Mehar Bhatia, Xiaofeng Zhang, Verena Rieser, Lisa Anne Hendricks, Sjoerd Van Steenkiste, Yash Goyal, Karolina Stanczak and Aishwarya Agrawal	20918
<i>Decoupled Proxy Alignment: Mitigating Language Prior Conflict for Multimodal Alignment in MLLMs</i>	
Chenkun Tan, Pengyu Wang, Shaojun Zhou, Botian Jiang, Zhaowei Li, Dong Zhang, Xinghao Wang, Yaqian Zhou and Xipeng Qiu	20954
<i>Riemannian Optimization for LoRA on the Stiefel Manifold</i>	
JuneYoung Park, Minjae Kang, Seongbae Lee, Haegang Lee, Seongwan Kim and Jaeho Lee	20971
<i>How Real Are Synthetic Therapy Conversations? Evaluating Fidelity in Prolonged Exposure Dialogues</i>	
Suhas BN, Dominik O. Mattioli, Andrew M. Sherrill, Rosa I. Arriaga, Christopher Wiese and Saeed Abdullah	20986

<i>Large Language Models for Controllable Multi-property Multi-objective Molecule Optimization</i> Vishal Dey, Xiao Hu and Xia Ning	20996
<i>Measuring Lexical Diversity of Synthetic Data Generated through Fine-Grained Persona Prompting</i> Gauri Kambhatla, Chantal Shaib and Venkata S Govindarajan	21024
<i>Beyond Function-Level Search: Repository-Aware Dual-Encoder Code Retrieval with Adversarial Verification</i> Aofan Liu, Song Shiyuan, Haoxuan li, Cehao Yang and Yiyan Qi	21034
<i>Watermark under Fire: A Robustness Evaluation of LLM Watermarking</i> Jiacheng Liang, Zian Wang, Spencer Hong, Shouling Ji and Ting Wang	21050
<i>PEPE: Long-context Extension for Large Language Models via Periodic Extrapolation Positional Encodings</i> Jikun Hu, Dongsheng Guo, Yuli Liu, Qingyao Ai, Lixuan Wang, Xuebing Sun, Qilei Zhang, Quan Zhou and Cheng Luo	21075
<i>Beyond Self-Reports: Multi-Observer Agents for Personality Assessment in Large Language Models</i> Yin Jou Huang and Rafik Hadfi	21086
<i>Controlled Retrieval-augmented Context Evaluation for Long-form RAG</i> Jia-Huei Ju, Suzan Verberne, Maarten de Rijke and Andrew Yates	21102
<i>Humanity's Last Code Exam: Can Advanced LLMs Conquer Human's Hardest Code Competition?</i> Xiangyang Li, Xiaopeng Li, Kuicai Dong, Zhangquanhu, Rongju Ruan, Xinyi Dai, Yasheng Wang and Ruiming Tang	21122
<i>False Friends Are Not Foes: Investigating Vocabulary Overlap in Multilingual Language Models</i> Julie Kallini, Dan Jurafsky, Christopher Potts and Martijn Bartelds	21138
<i>Rule-Guided Extraction: A Hierarchical Rule Optimization Framework for Document-Level Event Argument Extraction</i> Yue Zuo, Yuxiao Fei, Wanting Ning, Jiayi Huang, Yubo Feng and Lishuang Li	21155
<i>SOPL: A Sequential Optimal Learning Approach to Automated Prompt Engineering in Large Language Models</i> Shuyang Wang, Somayeh Moazeni and Diego Klabjan	21172
<i>CLIP-UP: A Simple and Efficient Mixture-of-Experts CLIP Training Recipe with Sparse Upcycling</i> Xinze Wang, Chen Chen, Yinfei Yang, Hong-You Chen, Bowen Zhang, Aditya Pal, Xiangxin Zhu and Xianzhi Du	21186
<i>A Category-Theoretic Approach to Neural-Symbolic Task Planning with Bidirectional Search</i> Shuhui Qu, Jie Wang and Kincho Law	21201
<i>HEAL: An Empirical Study on Hallucinations in Embodied Agents Driven by Large Language Models</i> Trishna Chakraborty, Udit Ghosh, Xiaopan Zhang, Fahim Faisal Niloy, Yue Dong, Jiachen Li, Amit Roy-Chowdhury and Chengyu Song	21226
<i>Can LLMs Judge Debates? Evaluating Non-Linear Reasoning via Argumentation Theory Semantics</i> Reza Sanayei, Srdjan Vesic, Eduardo Blanco and Mihai Surdeanu	21244
<i>How Jailbreak Defenses Work and Ensemble? A Mechanistic Investigation</i> Zhuohan Long, Siyuan Wang, Shujun Liu and Yuhang Lai	21263

<i>Visual Self-Refinement for Autoregressive Models</i>	
Jiamian Wang, Ziqi Zhou, Chaithanya Kumar Mummadi, Sohail Dianat, Majid Rabbani, Raghuv veer Rao, Chen Qiu and Zhiqiang Tao	21291
<i>Retrieval-Augmented Language Models are Mimetic Theorem Provers</i>	
Wenjie Yang, Ruiyuan Huang, Jiaying Guo, Zicheng Lyu, Tongshan Xu, Shengzhong Zhang, Lun Du, Da Zheng and Zengfeng Huang	21301
<i>LORE: Continual Logit Rewriting Fosters Faithful Generation</i>	
Charles Yu, Qingyun Wang, Yuting Hu, Jinjun Xiong and Heng Ji	21314
<i>PRINCIPLES: Synthetic Strategy Memory for Proactive Dialogue Agents</i>	
Namyong Kim, Kai Tzu-iunn Ong, Yeonjun Hwang, Minseok Kang, Iiseo Jihn, Gayoung Kim, Minju Kim and Jinyoung Yeo	21329
<i>SLM-Bench: A Comprehensive Benchmark of Small Language Models on Environmental Impacts</i>	
Nghiem Thanh Pham, Tung Kieu, Duc Manh Nguyen, Son Ha Xuan, Nghia Duong-Trung and Danh Le-Phuoc	21369
<i>A Decoupled Multi-Agent Framework for Complex Text Style Transfer</i>	
Lingxi Zhang, Yu-Neng Chuang, Guanchu Wang, Ruixiang Tang, Xuanting Cai, Rajesh Shenoy and Xia Hu	21393
<i>Mamba Drafters for Speculative Decoding</i>	
Daewon Choi, Seunghyuk Oh, Saket Dingliwal, Jihoon Tack, Kyuyoung Kim, Woomin Song, Seojin Kim, Insu Han, Jinwoo Shin, Aram Galstyan, Shubham Katiyar and Sravan Babu Bodapati 21404	
<i>LongLLaVA: Scaling Multi-modal LLMs to 1000 Images Efficiently via a Hybrid Architecture</i>	
Xidong Wang, Dingjie Song, Shunian Chen, Junying Chen, Zhenyang Cai, Chen Zhang, Lichao Sun and Benyou Wang	21419
<i>Think Clearly: Improving Reasoning via Redundant Token Pruning</i>	
Daewon Choi, Jimin Lee, Jihoon Tack, Woomin Song, Saket Dingliwal, Sai Muralidhar Jayanthi, Bhavana Ganesh, Jinwoo Shin, Aram Galstyan and Sravan Babu Bodapati	21437
<i>A Systematic Survey of Claim Verification: Corpora, Systems, and Case Studies</i>	
Zhaxi Zerong, Chenxi LI, Xinyi Liu, Ju-hui Chen and Fei Xia	21452
<i>Automated Creativity Evaluation for Large Language Models: A Reference-Based Approach</i>	
Ruizhe Li, Chiwei Zhu, Benfeng Xu, Xiaorui Wang and Zhendong Mao	21475
<i>LangProBe: a Language Program Benchmark</i>	
Shangyin Tan, Lakshya A Agrawal, Arnav Singhvi, Liheng Lai, Michael J Ryan, Dan Klein, Omar Khattab, Koushik Sen and Matei Zaharia	21489
<i>Exploring and Detecting Self-disclosure in Multi-modal posts on Chinese Social Media</i>	
Jingbao Luo, Ming Liu, Aoli Huo, Fujing Hu, Gang Li and Wupeng Njust	21510
<i>MV-CLAM: Multi-View Molecular Interpretation with Cross-Modal Projection via Language Model</i>	
Sumin Ha, Jun Hyeong Kim, Yinhua Piao, Changyun Cho and Sun Kim	21528
<i>Mind the Style Gap: Meta-Evaluation of Style and Attribute Transfer Metrics</i>	
Amalie Brogaard Pauli, Isabelle Augenstein and Ira Assent	21550
<i>ExtremeAIGC: Benchmarking LMM Vulnerability to AI-Generated Extremist Content</i>	
Bhavik Chandna, Mariam Aboujenane and Usman Naseem	21565

<i>Data Augmentation for Maltese NLP using Transliterated and Machine Translated Arabic Data</i>	
Kurt Micallef, Nizar Habash and Claudia Borg	21580
<i>Do LLMs Align Human Values Regarding Social Biases? Judging and Explaining Social Biases with LLMs</i>	
Yang Liu and Chenhui Chu	21591
<i>CoEx – Co-evolving World-model and Exploration</i>	
Minsoo Kim and Seung-won Hwang	21629
<i>BrainLoc: Brain Signal-Based Object Detection with Multi-modal Alignment</i>	
Jiaqi Duan, Xiaoda Yang, Kaixuan Luan, Hongshun Qiu, Weicai Yan, Xueyi Zhang, Youliang Zhang, Zhaoyang Li, Donglin Huang, JunYu Lu, Ziyue Jiang and Xifeng Yang	21652
<i>PVTNL: Prompting Vision Transformers with Natural Language for Generalizable Person Re-identification</i>	
Wangning, Lei Xie, Sanglu Lu and Shiwei Gan	21663
<i>RingFormer: Rethinking Recurrent Transformer with Adaptive Level Signals</i>	
Jaemu Heo, Eldor Fozilov, Hyunmin Song and Taehwan Kim	21675
<i>TriSPrompt: A Hierarchical Soft Prompt Model for Multimodal Rumor Detection with Incomplete Modalities</i>	
Jiajun Chen, Yangyang Wu, Xiaoye Miao, Mengying Zhu and Meng Xi	21687
<i>Evaluating Uncertainty Quantification Methods in Argumentative Large Language Models</i>	
Kevin Zhou, Adam Dejl, Gabriel Freedman, Lihu Chen, Antonio Rago and Francesca Toni	21700
<i>CLAIMCHECK: How Grounded are LLM Critiques of Scientific Papers?</i>	
Jiefu Ou, William Gantt Walden, Kate Sanders, Zhengping Jiang, Kaiser Sun, Jeffrey Cheng, William Jurajy, Miriam Wanner, Shaobo Liang, Candice Morgan, Seunghoon Han, Weiqi Wang, Chandler May, Hannah Recknor, Daniel Khashabi and Benjamin Van Durme	21712
<i>From Noise to Clarity: Filtering Real and LLM-Generated Samples for Enhanced Intent Detection</i>	
Junbao Huang, Weizhen Li, Peijie Huang and Yuhong Xu	21736
<i>Improving Language Model Personas via Rationalization with Psychological Scaffolds</i>	
Brihi Joshi, Xiang Ren, Swabha Swayamdipta, Rik Koncel-Kedziorski and Tim Paek	21747
<i>KBM: Delineating Knowledge Boundary for Adaptive Retrieval in Large Language Models</i>	
Zhen Zhang, Xinyu Wang, Yong Jiang, Zile Qiao, Zhuo Chen, Guangyu Li, Feiteng Mu, Mengting Hu, Pengjun Xie and Fei Huang	21771
<i>TABARD: A Novel Benchmark for Tabular Anomaly Analysis, Reasoning and Detection</i>	
Manan Roy Choudhury, Anirudh Iyengar Kaniyar Narayana Iyengar, Shikhar Singh, Sugeeth Puranam and Vivek Gupta	21783
<i>Aspect-based Sentiment Analysis via Synthetic Image Generation</i>	
Ge Chen, Zhongqing Wang and Guodong Zhou	21818
<i>IntrEx: A Dataset for Modeling Engagement in Educational Conversations</i>	
Xingwei Tan, Mahathi Parvatham, Chiara Gambi and Gabriele Pergola	21830
<i>Bridging the Capability Gap: Joint Alignment Tuning for Harmonizing LLM-based Multi-Agent Systems</i>	
Minghang Zhu, Zhengliang Shi, Zhiwei Xu, Shiguang Wu, Lingjie Wang, Pengjie Ren, Zhaochun Ren and Zhumin Chen	21846
<i>Safety Through Reasoning: An Empirical Study of Reasoning Guardrail Models</i>	
Makesh Narsimhan Sreedhar, Traian Rebedea and Christopher Parisien	21862

<i>Context-Aware Reasoning On Parametric Knowledge for Inferring Causal Variables</i> Ivaxi Sheth, Sahar Abdelnabi and Mario Fritz	21881
<i>LoRE-Merging: Exploring Low-Rank Estimation For Large Language Model Merging</i> Zehua Liu, Han Wu, Yuxuan Yao, Xiaojin Fu, Ruifeng She, Xiongwei Han, Tao Zhong and Mingxuan Yuan.....	21919
<i>Benchmarking Foundation Models with Retrieval-Augmented Generation in Olympic-Level Physics Problem Solving</i> Shunfeng Zheng, Yudi Zhang, Meng Fang, Zihan Zhang, Zhitan Wu, Mykola Pechenizkiy and Ling Chen.....	21927
<i>FiRST: Finetuning Router-Selective Transformers for Input-Adaptive Latency Reduction</i> Akriti Jain, Saransh Sharma, Koyel Mukherjee and Soumyabrata Pal.....	21957
<i>PolitiSky24: U.S. Political Bluesky Dataset with User Stance Labels</i> Peyman Rostami, Vahid Rahimzadeh, Ali Adibi and Azadeh Shakery	21976
<i>From Ground Trust to Truth: Disparities in Offensive Language Judgments on Contemporary Korean Political Discourse</i> Seunguk Yu, JungMin Yun, Jinhee Jang and YoungBin Kim	21994
<i>Misalignment Attack on Text-to-Image Models via Text Embedding Optimization and Inversion</i> Zhijie Du, Daizong Liu and Pan Zhou.....	22015
<i>Domain Pre-training Impact on Representations</i> Cesar Gonzalez-Gutierrez and Ariadna Quattoni	22033
<i>KoACD: The First Korean Adolescent Dataset for Cognitive Distortion Analysis via Role-Switching Multi-LLM Negotiation</i> Jun Seo Kim and Hye Hyeon Kim	22050
<i>Refined Assessment for Translation Evaluation: Rethinking Machine Translation Evaluation in the Era of Human-Level Systems</i> Dmitry Popov, Vladislav Negodin, Ekaterina Enikeeva, Iana Matrosova, Nikolay Karpachev and Max Ryabinin	22079
<i>Pre-Storage Reasoning for Episodic Memory: Shifting Inference Burden to Memory for Personalized Dialogue</i> Sangyeop Kim, Yohan Lee, Sanghwa Kim, Hyunjong Kim and Sungzoon Cho.....	22096
<i>Temporal Consistency for LLM Reasoning Process Error Identification</i> Jiacheng Guo, Yue Wu, Jiahao Qiu, Kaixuan Huang, Xinzhe Juan, Ling Yang and Mengdi Wang	22114
<i>Quantifying Compositionality of Classic and State-of-the-Art Embeddings</i> Zhijin Guo, Chenhao Xue, Zhaozhen Xu, Hongbo Bo, Yuxuan Ye, Janet B. Pierrehumbert and Martha Lewis	22130
<i>Presumed Cultural Identity: How Names Shape LLM Responses</i> Siddhesh Milind Pawar, Arnav Arora, Lucie-Aimée Kaffee and Isabelle Augenstein	22147
<i>I-GUARD: Interpretability-Guided Parameter Optimization for Adversarial Defense</i> Mamta Mamta and Oana Cocarascu	22173
<i>DecoupledESC: Enhancing Emotional Support Generation via Strategy-Response Decoupled Preference Optimization</i> Chao Zhang, Xin Shi, Xueqiao Zhang, Yifan Zhu, Yi Yang and Yawei Luo	22189

<i>Local Normalization Distortion and the Thermodynamic Formalism of Decoding Strategies for Large Language Models</i>	
Tom Kempton and Stuart Burrell	22216
<i>BRIT: Bidirectional Retrieval over Unified Image-Text Graph</i>	
Ainulla Khan, Moyuru Yamada and Srinidhi Akella	22232
<i>ReTAG: Retrieval-Enhanced, Topic-Augmented Graph-Based Global Sensemaking</i>	
Boyoung Kim, Dosung Lee, Sumin An, Jinseong Jeong and Paul Hongsuck Seo	22249
<i>Capturing Latent Modal Association For Multimodal Entity Alignment</i>	
Yongquan Ji, Jingwei Cheng, Fu Zhang and Chenglong Lu	22278
<i>Explaining novel senses using definition generation with open language models</i>	
Mariia Fedorova, Andrey Kutuzov, Francesco Periti and Yves Scherrer	22294
<i>Can Code-Switched Texts Activate a Knowledge Switch in LLMs? A Case Study on English-Korean Code-Switching</i>	
Soyeon Kim, Hwiseo Kim, Chanjun Park, Jinyoung Yeo and Dongha Lee	22303
<i>Compositional Translation: A Novel LLM-based Approach for Low-resource Machine Translation</i>	
Armel Randy Zebaze, Benoît Sagot and Rachel Bawden	22328
<i>TopXGen: Topic-Diverse Parallel Data Generation for Low-Resource Machine Translation</i>	
Armel Randy Zebaze, Benoît Sagot and Rachel Bawden	22358
<i>Fast, Not Fancy: Rethinking G2P with Rich Data and Statistical Models</i>	
Mahta Fetrat Qharabagh, Zahra Dehghanian and Hamid R. Rabiee	22382
<i>Personalized open world plan generation for safety-critical human centered autonomous systems: A case study on Artificial Pancreas</i>	
Ayan Banerjee and Sandeep Gupta	22409
<i>CaMMT: Benchmarking Culturally Aware Multimodal Machine Translation</i>	
Emilio Villa-Cueva, Sholpan Bolatzhanova, Diana Turmakhan, Kareem Elzeky, Henok Biadglign Ademtew, Alham Fikri Aji, Vladimir Araujo, Israel Abebe Azime, Jinheon Baek, Frederico Belcavello, Fermin Cristobal, Jan Christian Blaise Cruz, Mary Dabre, Raj Dabre, Toqeer Ehsan, Naome A Etori, Fauzan Farooqui, Jiahui Geng, Guido Ivetta, Thanmay Jayakumar, Soyeong Jeong, Zheng Wei Lim, Aishik Mandal, Sofia Martinelli, Mihail Minkov Mihaylov, Daniil Orel, Aniket Pramanick, Sukannya Purkayastha, Israfel Salazar, Haiyue Song, Tiago Timponi Torrent, Debela Desalegn Yadeta, Injy Hamed, Atnafu Lambebo Tonja and Tamar Solorio	22423
<i>Training Text-to-Molecule Models with Context-Aware Tokenization</i>	
Seojin Kim, Hyeontae Song, Jaehyun Nam and Jinwoo Shin	22442
<i>Challenging the Evaluator: LLM Sycophancy Under User Rebuttal</i>	
Sung Won Kim and Daniel Khashabi	22461
<i>Perspective-driven Preference Optimization with Entropy Maximization for Diverse Argument Generation</i>	
Yilin Cao, Ruike Zhang, Penghui Wei, Qingchao Kong and Wenji Mao	22479
<i>Spoken Document Retrieval for an Unwritten Language: A Case Study on Gormati</i>	
Sanjay Booshanam, Kelly Chen, Ondrej Klejch, Thomas Reitmaier, Dani Kalarikalayil Raju, Electra Wallington, Nina Markl, Jennifer Pearson, Matt Jones, Simon Robinson and Peter Bell ..	22497

<i>M-Help: Using Social Media Data to Detect Mental Health Help-Seeking Signals</i> Msvpj Sathvik, Zuhair Hasan Shaik and Vivek Gupta	22510
<i>Brittle Minds, Fixable Activations: Understanding Belief Representations in Language Models</i> Matteo Bortoletto, Constantin Ruhdorfer, Lei Shi and Andreas Bulling	22521
<i>Golden Touchstone: A Comprehensive Bilingual Benchmark for Evaluating Financial Large Language Models</i> Xiaojun Wu, Junxi Liu, Huan-Yi Su, Zhouchi Lin, Yiyan Qi, Chengjin Xu, Jiajun Su, Jiajie Zhong, Fuwei Wang, Saizhuo Wang, Fengrui Hua, Jia Li and Jian Guo	22544
<i>Quantifying the Risks of LLM- and Tool-assisted Rephrasing to Linguistic Diversity</i> Mengying Wang and Andreas Spitz	22561
<i>NUMINA: A Natural Understanding Benchmark for Multi-dimensional Intelligence and Numerical Reasoning Abilities</i> Changyu Zeng, Yifan Wang, Zimu Wang, Wei Wang, Zhengni Yang, Muyi Bao, Jimin Xiao, Anh Nguyen and Yutao Yue	22575
<i>MoMentS: A Comprehensive Multimodal Benchmark for Theory of Mind</i> Emilio Villa-Cueva, S M Masrur Ahmed, Rendi Chevi, Jan Christian Blaise Cruz, Kareem Elzaky, Fermin Cristobal, Alham Fikri Aji, Skyler Wang, Rada Mihalcea and Tamar Solorio	22591
<i>Code Like Humans: A Multi-Agent Solution for Medical Coding</i> Andreas Geert Motzfeldt, Joakim Edin, Casper L. Christensen, Christian Hardmeier, Lars Maaløe and Anna Rogers	22612
<i>Can Out-of-Distribution Evaluations Uncover Reliance on Prediction Shortcuts? A Case Study in Question Answering</i> Michal Štefánik, Timothee Mickus, Michal Spiegel, Marek Kadlčík and Josef Kuchař	22628
<i>MEXA: Towards General Multimodal Reasoning with Dynamic Multi-Expert Aggregation</i> Shoubin Yu, Yue Zhang, Ziyang Wang, Jaehong Yoon and Mohit Bansal	22636
<i>Lifelong Knowledge Editing requires Better Regularization</i> Akshat Gupta, Phudish Prateepamornkul, Maochuan Lu, Ahmed Alaa, Thomas Hartvigsen and Gopala Anumanchipalli	22653
<i>Lost in Embeddings: Information Loss in Vision–Language Models</i> Wenyan Li, Raphael Tang, Chengzu Li, Caiqi Zhang, Ivan Vulić and Anders Søgaard	22676
<i>Assessing the Role of Data Quality in Training Bilingual Language Models</i> Skyler Seto, Maartje Ter Hoeve, Maureen de Seyssel and David Grangier	22694
<i>DORM: Preference Data Weights Optimization for Reward Modeling in LLM Alignment</i> Rongzhi Zhang, Chenwei Zhang, Xinyang Zhang, Liang Qiu, Haoming Jiang, Yuchen Zhuang, Qingru Zhang, Hyokun Yun, Xian Li, Bing Yin, Tuo Zhao and Chao Zhang	22721
<i>Enhancing Domain-Specific Encoder Models with LLM-Generated Data: How to Leverage Ontologies, and How to Do Without Them</i> Marc Felix Brinner, Tarek Al Mustafa and Sina Zarrieß	22740
<i>Aligning Dialogue Agents with Global Feedback via Large Language Model Multimodal Reward Decomposition</i> Dong Won Lee, Hae Won Park, Cynthia Breazeal and Louis-Philippe Morency	22755

<i>UrduFactCheck: An Agentic Fact-Checking Framework for Urdu with Evidence Boosting and Benchmarking</i>	
Sarfraz Ahmad, Hasan Iqbal, Momina Ahsan, Numaan Naeem, Muhammad Ahsan Riaz Khan, Arham Riaz, Muhammad Arslan Manzoor, Yuxia Wang and Preslav Nakov	22788
<i>Echoes of Agreement: Argument Driven Sycophancy in Large Language models</i>	
Avneet Kaur	22803
<i>Rethinking NLP for Chemistry: A Critical Look at the USPTO Benchmark</i>	
Derin Ozer, Nicolas Gutowski, Benoit Da Mota, Thomas Cauchy and Sylvain Lamprier	22813
<i>Investigating Dictionary Expansion for Video-based Sign Language Dictionaries</i>	
Aashaka Desai, Daniela Massiceti, Richard Ladner, Hal Daumé Iii, Danielle Bragg and Alex Xijie Lu	22826
<i>From Insight to Exploit: Leveraging LLM Collaboration for Adaptive Adversarial Text Generation</i>	
Najrin Sultana, Md Rafi Ur Rashid, Kang Gu and Shagufta Mehnaz	22842
<i>Beyond Contrastive Learning: Synthetic Data Enables List-wise Training with Multiple Levels of Relevance</i>	
Reza Esfandiarpour, George Zerveas, Ruochen Zhang, Macton Mgonzo, Carsten Eickhoff and Stephen Bach	22860
<i>Instability in Downstream Task Performance During LLM Pretraining</i>	
Yuto Nishida, Masaru Isonuma and Yusuke Oda	22883
<i>A Comparison of Independent and Joint Fine-tuning Strategies for Retrieval-Augmented Generation</i>	
Neal Gregory Lawton, Alfy Samuel, Anoop Kumar and Daben Liu	22896
<i>mrCAD: Multimodal Communication to Refine Computer-aided Designs</i>	
William P McCarthy, Saujas Vaduguru, Karl D.d. Willis, Justin Matejka, Judith E Fan, Daniel Fried and Yewen Pu	22905
<i>MOCHA: Are Code Language Models Robust Against Multi-Turn Malicious Coding Prompts?</i>	
Muntasir Wahed, Xiaona Zhou, Kiet A. Nguyen, Tianjiao Yu, Nirav Diwan, Gang Wang, Dilek Hakkani-Tür and Ismini Lourentzou	22922
<i>How Can Input Reformulation Improve Tool Usage Accuracy in a Complex Dynamic Environment? A Study on tau-bench</i>	
Venkatesh Mishra, Amir Saeidi, Satyam Raj, Mutsumi Nakamura, Gaowen Liu, Ali Payani, Jayanth Srinivasa and Chitta Baral	22949
<i>Evaluating Fairness in Large Vision-Language Models Across Diverse Demographic Attributes and Prompts</i>	
Xuyang Wu, Yuan Wang, Hsin-Tai Wu, Zhiqiang Tao and Yi Fang	22973
<i>VIBE: Can a VLM Read the Room?</i>	
Tania Chakraborty, Eylon Caplan and Dan Goldwasser	22992
<i>LoRATK: LoRA Once, Backdoor Everywhere in the Share-and-Play Ecosystem</i>	
Hongyi Liu, Shaochen Zhong, Xintong Sun, Minghao Tian, Mohsen Hariri, Zirui Liu, Ruixiang Tang, Zhimeng Jiang, Jiayi Yuan, Yu-Neng Chuang, Li Li, Soo-Hyun Choi, Rui Chen, Vipin Chaudhary and Xia Hu	23009
<i>Pearl: A Multimodal Culturally-Aware Arabic Instruction Dataset</i>	
Fakhraddin Alwajih, Samar M. Magdy, Abdellah El Mekki, Omer Nacar, Youssef Nafea, Safaa	

Taher Abdelfadil, Abdulfattah Mohammed Yahya, Hamzah Luqman, Nada Almarwani, Samah Aloufi, Baraah Qawasmeh, Houdaifa Atou, Serry Sibae, Hamzah A. Alsayadi, Walid Al-Dhabyani, Maged S. Al-shaibani, Aya El aatar, Nour Qandos, Rahaf Alhamouri, Samar Ahmad, Mohammed Anwar AL-Ghrawi, Aminetou Yacoub, Ruwa AbuHweidi, Vatimetou Mohamed Lemin, Reem Abdel-Salam, Ahlam Bashiti, Adel Ammar, Aisha Alansari, Ahmed Ashraf, Nora Alturayef, Alcides Alcoba Inciarte, AbdelRahim A. Elmadany, Mohamedou Cheikh Tourad, Ismail Berrada, Mustafa Jarrar, Shady Shehata and Muhammad Abdul-Mageed	23048
<i>Protein Large Language Models: A Comprehensive Survey</i>	
Yijia Xiao, Wanxia Zhao, Junkai Zhang, Yiqiao Jin, Han Zhang, Zhicheng Ren, Renliang Sun, Haixin Wang, Guancheng Wan, Pan Lu, Xiao Luo, Yu Zhang, James Zou, Yizhou Sun and Wei Wang	23080
<i>MAKIEval: A Multilingual Automatic WiKidata-based Framework for Cultural Awareness Evaluation for LLMs</i>	
Raoyuan Zhao, Beiduo Chen, Barbara Plank and Michael A. Hedderich	23104
<i>Looking Beyond the Pixels: Evaluating Visual Metaphor Understanding in VLMs</i>	
Manishit Kundu, Sumit Shekhar and Pushpak Bhattacharyya	23137
<i>AGENTVIGIL: Automatic Black-Box Red-teaming for Indirect Prompt Injection against LLM Agents</i>	
Zhun Wang, Vincent Siu, Zhe Ye, Tianneng Shi, Yuzhou Nie, Xuandong Zhao, Chenguang Wang, Wenbo Guo and Dawn Song	23159
<i>Improving LLM-as-a-Judge Inference with the Judgment Distribution</i>	
Victor Wang, Michael JQ Zhang and Eunsol Choi	23173
<i>Learning Is Not A Race: Improving Retrieval in Language Models via Equal Learning</i>	
Wanqian Yang, Aahlad Manas Puli and Rajesh Ranganath	23200
<i>The Prompt Makes the Person(a): A Systematic Evaluation of Sociodemographic Persona Prompting for Large Language Models</i>	
Marlene Lutz, Indira Sen, Georg Ahnert, Elisa Rogers and Markus Strohmaier	23212
<i>Spiral of Silence in Large Language Model Agents</i>	
Mingze Zhong, Meng Fang, Zijing Shi, Yuxuan Huang, Shunfeng Zheng, Yali Du, Ling Chen and Jun Wang	23238
<i>Do We Know What LLMs Don't Know? A Study of Consistency in Knowledge Probing</i>	
Raoyuan Zhao, Abdullatif Köksal, Ali Modarressi, Michael A. Hedderich and Hinrich Schuetze	23254
<i>Context Length Alone Hurts LLM Performance Despite Perfect Retrieval</i>	
Yufeng Du, Minyang Tian, Srikanth Ronanki, Subendhu Rongali, Sravan Babu Bodapati, Aram Galstyan, Azton Wells, Roy Schwartz, Eliu A Huerta and Hao Peng	23281
<i>DebUnc: Improving Large Language Model Agent Communication With Uncertainty Metrics</i>	
Luke Yoffe, Alfonso Amayuelas and William Yang Wang	23299
<i>ProcVQA: Benchmarking the Effects of Structural Properties in Mined Process Visualizations on Vision–Language Model Performance</i>	
Kazi Tasnim Zinat, Saad Mohammad Abrar, Shoumik Saha, Sharmila Duppala, Saimadhav Naga Sakhamuri and Zhicheng Liu	23316
<i>Probing Political Ideology in Large Language Models: How Latent Political Representations Generalize Across Tasks</i>	
Tianyi Zhang	23349

<i>Understanding GUI Agent Localization Biases through Logit Sharpness</i> Xingjian Tao, Yiwei Wang, Yujun Cai, Zhicheng Yang and Jing Tang	23361
<i>The Language of Interoception: Examining Embodiment and Emotion Through a Corpus of Body Part Mentions</i> Sophie Wu, Jan Philip Wahle and Saif M. Mohammad	23375
<i>HomoGraphAdapter: A Homogeneous Graph Neural Network as an Effective Adapter for Vision-Language Models</i> Chuan He, Zhuozhao Li, Song Guo, Xiaocheng Lu and Jinxiang Lai	23400
<i>No Black Boxes: Interpretable and Interactable Predictive Healthcare with Knowledge-Enhanced Agentic Causal Discovery</i> Xiaoxue Han, Pengfei Hu, Chang Lu, Jun-En Ding, Feng Liu and Yue Ning	23415
<i>PROOD: A Simple LLM Out-of-Distribution Guardrail Leveraging Response Semantics</i> Joshua Tint	23428
<i>ICL-Bandit: Relevance Labeling in Advertisement Recommendation Systems via LLM</i> Lu Wang, Chiming Duan, Pu Zhao, Fangkai Yang, Yong Shi, Xuefeng Luo, Bingjing Xu, Weiwei Deng, Qingwei Lin and Dongmei Zhang	23439
<i>Intent-aware Schema Generation and Refinement for Literature Review Tables</i> Vishakh Padmakumar, Joseph Chee Chang, Kyle Lo, Doug Downey and Aakanksha Naik .	23450
<i>NLP Needs Diversity outside of 'Diversity'</i> Joshua Tint	23473
<i>Anatomy of a Feeling: Narrating Embodied Emotions via Large Vision-Language Models</i> Mohammad Saim, Phan Anh Duong, Cat Luong, Aniket Bhanderi and Tianyu Jiang	23480
<i>Towards Universal Debiasing for Language Models-based Tabular Data Generation</i> Tianchun Li, Tianci Liu, Xingchen Wang, Rongzhe Wei, Pan Li, Lu Su and Jing Gao	23496
<i>Beyond Linear Steering: Unified Multi-Attribute Control for Language Models</i> Narmeen Fatimah Oozeer, Luke Marks, Fazl Barez and Amir Abdullah	23513
<i>Unequal Scientific Recognition in the Age of LLMs</i> Yixuan Liu, Abel Elekes, Jianglin Lu, Rodrigo Dorantes-Gilardi and Albert-Laszlo Barabasi	23558
<i>Zero-Shot Fine-Grained Image Classification Using Large Vision-Language Models</i> Md. Atabuzzaman, Andrew Zhang and Chris Thomas	23569
<i>Using tournaments to calculate AUROC for zero-shot classification with LLMs</i> WonJin Yoon, Ian Bulovic and Timothy A. Miller	23583
<i>Exploration-Driven Reinforcement Learning for Expert Routing Improvement in Mixture-of-Experts Language Models</i> Gyunyeop Kim and Sangwoo Kang	23592
<i>D2CS - Documents Graph Clustering using LLM supervision</i> Yoel Ashkenazi, Etzion Harari, Regev Yehezkel Imra, Naphtali Abudarham, Dekel Cohen and Yoram Louzoun	23606
<i>GeoChain: Multimodal Chain-of-Thought for Geographic Reasoning</i> Sahiti Yerramilli, Nilay Pande, Rynaa Grover and Jayant Sravan Tamarapalli	23624

<i>SteerVLM: Robust Model Control through Lightweight Activation Steering for Vision Language Models</i> Anushka Sivakumar, Andrew Zhang, Zaber Ibn Abdul Hakim and Chris Thomas	23640
<i>FractalLLM: Lossless Self-Speculative Decoding with Layer Embedded Self-Compression</i> Juhyeong Kim, Sangyeon Yu, Gyunyeop Kim and Sangwoo Kang	23666
<i>Saten: Sparse Augmented Tensor Networks for Post-Training Compression of Large Language Models</i> Ryan Solgi, Kai Zhen, Rupak Vignesh Swaminathan, Nathan Susanj, Athanasios Mouchtaris, Siegfried Kunzmann and Zheng Zhang	23674
<i>Third-Person Appraisal Agent: Simulating Human Emotional Reasoning in Text with Large Language Models</i> Simin Hong, Jun Sun and Hongyang Chen	23684
<i>Source-primed Multi-turn Conversation Helps Large Language Models Translate Documents</i> Hanxu Hu, Jannis Vamvas and Rico Sennrich	23702
<i>Mitigating Spurious Correlations via Counterfactual Contrastive Learning</i> Fengxiang Cheng, Chuan Zhou, Xiang Li, Alina Leidinger, Haoxuan Li, Mingming Gong, Fen- rong Liu and Robert Van Rooij	23713
<i>The RAG Paradox: A Black-Box Attack Exploiting Unintentional Vulnerabilities in Retrieval-Augmented Generation Systems</i> Chanwoo Choi, Jinsoo Kim, Sukmin Cho, Soyeong Jeong and Buru Chang	23723
<i>Guiding Large Language Models for Biomedical Entity Linking via Restrictive and Contrastive Decoding</i> Zhenxi Lin, Ziheng Zhang, Jian Wu, Yefeng Zheng and Xian Wu	23745
<i>Cut the Deadwood Out: Backdoor Purification via Guided Module Substitution</i> Yao Tong, Weijun Li, Xuanli He, Haolan Zhan and Qionгкаi Xu	23760
<i>RepoDebug: Repository-Level Multi-Task and Multi-Language Debugging Evaluation of Large Language Models</i> Jingjing Liu, Zeming Liu, Zihao Cheng, Mengliang He, Xiaoming Shi, Yuhang Guo, Xiangrong Zhu, Yuanfang Guo, Yunhong Wang and Haifeng Wang	23784
<i>FaStFact: Faster, Stronger Long-Form Factuality Evaluations in LLMs</i> Yingjia Wan, Haochen Tan, Xiao Zhu, Xinyu Zhou, Zhiwei Li, Qingsong Lv, Changxuan Sun, Jiaqi Zeng, Yi Xu, Jianqiao Lu, Yinhong Liu and Zhijiang Guo	23814
<i>PropXplain: Can LLMs Enable Explainable Propaganda Detection?</i> Maram Hasanain, Md Arid Hasan, Mohamed Bayan Kmainasi, Elisa Sartori, Ali Ezzat Shahroor, Giovanni Da San Martino and Firoj Alam	23855
<i>EoT: Evolution of Thoughts for Complex Reasoning Tasks</i> Qin Hua, Jiaqi Sun, Shiyu Qian, Dingyu Yang, Jian Cao and Guangtao Xue	23864
<i>Reveal and Release: Iterative LLM Unlearning with Self-generated Data</i> Linxi Xie, Xin Teng, Shichang Ke, Hongyi Wen and Shenji Wan	23887
<i>An Evaluation Resource for Grounding Translation Errors</i> Sujin Chen, Kang Wang, Zixuan Zhou, Xiangyu Duan, Wanqun Zhang, Hao Yang, Jinsong Su and Min Zhang	23900
<i>Enhancing Time Awareness in Generative Recommendation</i> Sunkyoung Lee, Seongmin Park, Jonghyo Kim, Mincheol Yoon and Jongwuk Lee	23917

<i>Adaptive LLM Routing under Budget Constraints</i>	
Pranoy Panda, Raghav Magazine, Chaitanya Devaguptapu, Sho Takemori and Vishal Sharma	
23934	
<i>Promptception: How Sensitive Are Large Multimodal Models to Prompts?</i>	
Mohamed Insaf Ismithdeen, Muhammad Uzair Khattak and Salman Khan	23950
<i>Can Federated Learning Safeguard Private Data in LLM Training? Vulnerabilities, Attacks, and Defense Evaluation</i>	
Wenkai Guo, Xuefeng Liu, Haolin Wang, Jianwei Niu, Shaojie Tang and Jing Yuan	23986
<i>Runaway is Ashamed, But Helpful: On the Early-Exit Behavior of Large Language Model-based Agents in Embodied Environments</i>	
Qingyu Lu, Liang Ding, Siyi Cao, Xuebo Liu, Kanjian Zhang, Jinxia Zhang and Dacheng Tao	24014
<i>AutoMIR: Effective Zero-Shot Medical Information Retrieval without Relevance Labels</i>	
Lei Li, Xiangxu Zhang, Xiao Zhou and Zheng Liu	24028
<i>RG-VQA: Leveraging Retriever-Generator Pipelines for Knowledge Intensive Visual Question Answering</i>	
Settaluri Lakshmi Sravanthi, Pulkit Agarwal, Debjyoti Mondal, Rituraj Singh, Subhadarshi Panda, Ankit Mishra, Kiran Pradeep, Srihari K B, Godawari Sudhakar Rao and Pushpak Bhattacharyya	24048
<i>Enhancing RAG Efficiency with Adaptive Context Compression</i>	
Shuyu Guo, Shuo Zhang and Zhaochun Ren	24061
<i>Revealing the impact of synthetic native samples and multi-tasking strategies in Hindi-English code-mixed humour and sarcasm detection</i>	
Debajyoti Mazumder, Aakash Kumar and Jasabanta Patro	24077
<i>CogAtom: From Cognitive Atoms to Olympiad-level Mathematical Reasoning in Large Language Models</i>	
Zhuofan Chen, Jiyuan He, Yichi Zhang, Xing Hu, Haoxing Wen, Jun Bai and Wenge Rong	24108
<i>Efficient Latent Semantic Clustering for Scaling Test-Time Computation of LLMs</i>	
Sungjae Lee, Hoyoung Kim, Jeongyeon Hwang, Eunhyeok Park and Jungseul Ok	24126
<i>BannerBench: Benchmarking Vision Language Models for Multi-Ad Selection with Human Preferences</i>	
Hiroto Otake, Peinan Zhang, Yusuke Sakai, Masato Mita, Hiroki Ouchi and Taro Watanabe	24145
<i>DeKeyNLU: Enhancing Natural Language to SQL Generation through Task Decomposition and Keyword Extraction</i>	
Jian Chen, Zhenyan Chen, Xuming Hu, Peilin Zhou, Yining Hua, Han Fang, Cissy Hing Yee Choy, Xinmei Ke, Jingfeng Luo and Zixuan Yuan	24160
<i>Facilitating Cross-lingual Transfer of Empathy through Language-independent Latent Diffusion: A Case Study in Chinese</i>	
Junlin Li, Peng Bo and Yu-Yin Hsu	24177
<i>Evaluating Compound AI Systems through Behaviors, Not Benchmarks</i>	
Pranav Bhagat, K N Ajay Shastry, Pranoy Panda and Chaitanya Devaguptapu	24193
<i>SciCompanion: Graph-Grounded Reasoning for Structured Evaluation of Scientific Arguments</i>	
Joshua Alan Flashner, Adithya Kulkarni and Dawei Zhou	24223

<i>From Generation to Detection: A Multimodal Multi-Task Dataset for Benchmarking Health Misinformation</i>	
Zhihao Zhang, Yiran Zhang, Xiyue Zhou, Liting Huang, Imran Razzak, Preslav Nakov and Usman Naseem	24245
<i>Estimating Machine Translation Difficulty</i>	
Lorenzo Proietti, Stefano Perrella, Vilém Zouhar, Roberto Navigli and Tom Kocmi	24261
<i>TIU-Bench: A Benchmark for Evaluating Large Multimodal Models on Text-rich Image Understanding</i>	
Kun Zhang, Liqiang Niu, Zhen Cao, Fandong Meng and Jie Zhou	24286
<i>Breaking Token Into Concepts: Exploring Extreme Compression in Token Representation Via Compositional Shared Semantics</i>	
Kavin R V and Pawan Goyal	24296
<i>ExeSQL: Self-Taught Text-to-SQL Models with Execution-Driven Bootstrapping for SQL Dialects</i>	
Jipeng Zhang, Haolin Yang, Kehao Miao, Ruiyuan Zhang, Renjie Pi, Jiahui Gao and Xiaofang Zhou	24305
<i>Under the Shadow of Babel: How Language Shapes Reasoning in LLMs</i>	
Chenxi Wang, Yixuan Zhang, Lang Gao, Zixiang Xu, Zirui Song, Yanbo Wang and Xiuying Chen	24327
<i>Think Right, Not More: Test-Time Scaling for Numerical Claim Verification</i>	
Primakov Chungkham, Venkatesh V, Vinay Setty and Avishek Anand	24345
<i>Nexus: Adaptive Upcycling to Efficiently Pretrain Mixture of Experts</i>	
Nikolas Gritsch, Qizhen Zhang, Acyr Locatelli, Sara Hooker and Ahmet Üstün	24364
<i>Exploring Context Strategies in LLMs for Discourse-Aware Machine Translation</i>	
Ritvik Choudhary, Rem Hida, Masaki Hamada, Hayato Futami and Toshiyuki Sekiya	24382
<i>Insights into using temporal coordinated behaviour to explore connections between social media posts and influence</i>	
Elisa Sartori, Serena Tardelli, Maurizio Tesconi, Mauro Conti, Alessandro Galeazzi, Stefano Cresci and Giovanni Da San Martino	24392
<i>SpecCoT: Accelerating Chain-of-Thought Reasoning through Speculative Exploration</i>	
Junhan Shi, Yijia Zhu, Zhenning Shi, Dan Zhao, Qing Li and Yong Jiang	24405
<i>A Similarity Measure for Comparing Conversational Dynamics</i>	
Sang Min Jung, Kaixiang Zhang and Cristian Danescu-Niculescu-Mizil	24416
<i>AgentDrug: Utilizing Large Language Models in an Agentic Workflow for Zero-Shot Molecular Optimization</i>	
Le Huy Khiem, Ting Hua and Nitesh V Chawla	24448
<i>Improving Preference Alignment of LLM with Inference-Free Self-Refinement</i>	
Fukun Ma, Kaibin Tian, Jieting Xue, Xiaoyi Wang, Ye Ma, Quan Chen, Peng Jiang and Lijie Wen	24459
<i>Guaranteed Guess: A Language Modeling Approach for CISC-to-RISC Transpilation with Testing Guarantees</i>	
Ahmed Heakl, Sarim Hashmi, Chaimaa Abi, Celine Lee and Abdulrahman Mahmoud	24474
<i>StructuThink: Reasoning with Task Transition Knowledge for Autonomous LLM-Based Agents</i>	
Haiyu Zhao, Zhenyu Guo, Chunhong Zhang, Ziyu Zhou and Zheng Hu	24489

<i>Leveraging Unpaired Feedback for Long-Term LLM-based Recommendation Tuning</i> Jizhi Zhang, Chongming Gao, Wentao Shi, Xin Chen, Jingang Wang, Xunliang Cai and Fuli Feng	24507
<i>Investigating Multi-layer Representations for Dense Passage Retrieval</i> Zhongbin Xie and Thomas Lukasiewicz	24522
<i>KELE: Residual Knowledge Erasure for Enhanced Multi-hop Reasoning in Knowledge Editing</i> Mengqi Zhang, Bowen Fang, Qiang Liu, Xiaotian Ye, Shu Wu, Pengjie Ren, Zhumin Chen and Liang Wang	24537
<i>Dissecting Persona-Driven Reasoning in Language Models via Activation Patching</i> Ansh Poonia and Maeghal Jain	24553
<i>PUER: Boosting Few-shot Positive-Unlabeled Entity Resolution with Reinforcement Learning</i> Yaoshu Wang, Mengyi Yan and Wei Wang	24567
<i>Toward the Automatic Detection of Word Meaning Negotiation Indicators in Conversation</i> Aina Garí Soler, Matthieu Labeau and Chloé Clavel	24580
<i>Forget the Unneeded: Backdooring Large Language Models via Contrastive-enhanced Machine Un-learning</i> Shiji Yang, Shu Zhao, Congyao Mei, Zhen Yang, Jie Chen, Fulan Qian, Zhen Duan and Yanping Zhang	24597
<i>Equipping Retrieval-Augmented Large Language Models with Document Structure Awareness</i> Lingnan Xu, Chong Feng, Kaiyuan Zhang, Liu Zhengyong, Wenqiang Xu and Fanqing Meng	24608
<i>QEVA: A Reference-Free Evaluation Metric for Narrative Video Summarization with Multimodal Question Answering</i> Woojun Jung and Junyeong Kim	24632
<i>Thinking Before You Speak: A Proactive Test-time Scaling Approach</i> Cong Liu, Wenchang Chai, Hejun Wu, Yan Pan, Pengxu Wei and Liang Lin	24643
<i>Do Before You Judge: Self-Reference as a Pathway to Better LLM Evaluation</i> Wei-Hsiang Lin, Sheng-Lun Wei, Hen-Hsen Huang and Hsin-Hsi Chen	24651
<i>Beyond Content: How Grammatical Gender Shapes Visual Representation in Text-to-Image Models</i> Muhammed Saeed, Shaina Raza, Ashmal Vayani, Muhammad Abdul-Mageed, Ali Emami and Shady Shehata	24673
<i>ToolHaystack: Stress-Testing Tool-Augmented Language Models in Realistic Long-Term Interactions</i> Beong-woo Kwak, Minju Kim, Dongha Lim, Hyungjoo Chae, Dongjin Kang, Sunghwan Kim, Dongil Yang and Jinyoung Yeo	24696
<i>GraphCheck: Multipath Fact-Checking with Entity-Relationship Graphs</i> Hyeon Jeon and Jay-Yoon Lee	24728
<i>FLAMES: Improving LLM Math Reasoning via a Fine-Grained Analysis of the Data Synthesis Pipeline</i> Parker Seegmiller, Kartik Mehta, Soumya Saha, Chenyang Tao, Shereen Oraby, Arpit Gupta, Tagyoung Chung, Mohit Bansal and Nanyun Peng	24746
<i>POW: Political Overton Windows of Large Language Models</i> Leif Azzopardi and Yashar Moshfeghi	24767

<i>Columbo: Expanding Abbreviated Column Names for Tabular Data Using Large Language Models</i> Ting Cai, Stephen Sheen and AnHai Doan	24774
<i>RTTC: Reward-Guided Collaborative Test-Time Compute</i> Juan Pablo Munoz and Jinjie Yuan	24793
<i>AMANDA: Agentic Medical Knowledge Augmentation for Data-Efficient Medical Visual Question Answering</i> Ziqing Wang, Chengsheng Mao, Xiaole Wen, Yuan Luo and Kaize Ding	24810
<i>Mixed Signals: Decoding VLMs' Reasoning and Underlying Bias in Vision-Language Conflict</i> Pouya Pezeshkpour, Moin Aminnaseri and Estevam Hruschka	24833
<i>Mitigating Hallucination in Large Vision-Language Models through Aligning Attention Distribution to Information Flow</i> Jianfei Zhao, Feng Zhang, Xin Sun and Chong Feng	24849
<i>OptiSeq: Ordering Examples On-The-Fly for In-Context Learning</i> Rahul Atul Bhope, Praveen Venkateswaran, K. R. Jayaram, Vatche Isahagian, Vinod Muthusamy and Nalini Venkatasubramanian	24864
<i>Dependency Parsing-Based Syntactic Enhancement of Relation Extraction in Scientific Texts</i> Devvrat Joshi and Islem Rekik	24888
<i>DIPLOMA: Efficient Adaptation of Instructed LLMs to Low-Resource Languages via Post-Training Delta Merging</i> Ixak Sarasua Antero, Ander Corral and Xabier Saralegi	24898
<i>Reliability Crisis of Reference-free Metrics for Grammatical Error Correction</i> Takumi Goto, Yusuke Sakai and Taro Watanabe	24913
<i>Who Speaks Matters: Analysing the Influence of the Speaker's Linguistic Identity on Hate Classification</i> Ananya Malik, Kartik Sharma, Shaily Bhatt and Lynnette Hui Xian Ng	24927
<i>Are LLMs Empathetic to All? Investigating the Influence of Multi-Demographic Personas on a Model's Empathy</i> Ananya Malik, Nazanin Sabri, Melissa M. Karnaze and Mai ElSherief	24938
<i>Active Learning for Multidialectal Arabic POS Tagging</i> Diyam Akra, Mohammed Khalilia and Mustafa Jarrar	24960
<i>Embedding-Free RAG</i> Jessica Maghakian, Raunak Sinha, Max Schettewi and Gunkirat Kaur	24974
<i>Rating Roulette: Self-Inconsistency in LLM-As-A-Judge Frameworks</i> Rajarshi Haldar and Julia Hockenmaier	24986
<i>Quantifying Uncertainty in Natural Language Explanations of Large Language Models for Question Answering</i> Yangyi Li and Mengdi Huai	25005
<i>Real-World Summarization: When Evaluation Reaches Its Limits</i> Patrícia Schmidtová, Ondrej Dusek and Saad Mahamood	25014
<i>Open-DeBias: Toward Mitigating Open-Set Bias in Language Models</i> Arti Rani, Shweta Singh, Nihar Ranjan Sahoo and Gaurav Kumar Nayak	25027

<i>SACL: Understanding and Combating Textual Bias in Code Retrieval with Semantic-Augmented Reranking and Localization</i>	
Dhruv Gupta, Gayathri Ganesh Lakshmy and Yiqing Xie	25052
<i>Jailbreak Distillation: Renewable Safety Benchmarking</i>	
Jingyu Zhang, Ahmed Elgohary, Xiawei Wang, A S M Iftekhar, Ahmed Magooda, Benjamin Van Durme, Daniel Khashabi and Kyle Jackson	25066
<i>Uncertainty-Aware Answer Selection for Improved Reasoning in Multi-LLM Systems</i>	
Aakriti Agrawal, Rohith Aralikatti, Anirudh Satheesh, Souradip Chakraborty, Amrit Singh Bedi and Furong Huang	25090
<i>GreekBarBench: A Challenging Benchmark for Free-Text Legal Reasoning and Citations</i>	
Odysseas S. Chlapanis, Dimitris Galanis, Nikolaos Aletras and Ion Androutsopoulos	25099
<i>Pi-SQL: Enhancing Text-to-SQL with Fine-Grained Guidance from Pivot Programming Languages</i>	
Yongdong Chi, Hanqing Wang, Yun Chen, Yan Yang, Jian Yang, Zonghan Yang, Xiao Yan and Guanhua Chen	25120
<i>RAC: Efficient LLM Factuality Correction with Retrieval Augmentation</i>	
Changmao Li and Jeffrey Flanigan	25145
<i>Does It Run and Is That Enough? Revisiting Text-to-Chart Generation with a Multi-Agent Approach</i>	
James Ford and Anthony Rios	25160
<i>GeLoRA: Geometric Adaptive Ranks For Efficient LoRA Fine-tuning</i>	
Abdessalam Ed-dib, Zhanibek Datbayev and Amine M. Aboussalah	25174
<i>Uncovering Scaling Laws for Large Language Models via Inverse Problems</i>	
Arun Verma, Zhaoxuan Wu, Zijian Zhou, Xiaoqiang Lin, Zhiliang Chen, Rachael Hwee Ling Sim, Rui Qiao, Jingtang Wang, Nhung Bui, Xinyuan Niu, Wenyang Hu, Gregory Kang Ruey Lau, Zi-Yu Khoo, Zitong Zhao, Xinyi Xu, Apivich Hemachandra, See-Kiong Ng and Bryan Kian Hsiang Low	25197
<i>UIPE: Enhancing LLM Unlearning by Removing Knowledge Related to Forgetting Targets</i>	
Wenyu Wang, Mengqi Zhang, Xiaotian Ye, Zhaochun Ren, Pengjie Ren and Zhumin Chen	25212
<i>FicSim: A Dataset for Multi-Faceted Semantic Similarity in Long-Form Fiction</i>	
Natasha Johnson, Amanda Bertsch, Maria-Emil Deal and Emma Strubell	25228
<i>Masked Diffusion Captioning for Visual Feature Learning</i>	
Chao Feng, Zihao Wei and Andrew Owens	25247
<i>Diverse Multi-tool Aggregation with Large Language Models for Enhanced Math Reasoning</i>	
Bohan Yao and Vikas Yadav	25264
<i>Enhancing Goal-oriented Proactive Dialogue Systems via Dynamic Multi-dimensional Consistency Optimization</i>	
Didi Zhang, Yaxin Fan, Peifeng Li and Qiaoming Zhu	25283
<i>Injecting Domain-Specific Knowledge into Large Language Models: A Comprehensive Survey</i>	
Zirui Song, Bin Yan, Yuhan Liu, Miao Fang, Mingzhe Li, Rui Yan and Xiuying Chen	25297
<i>Who's the Author? How Explanations Impact User Reliance in AI-Assisted Authorship Attribution</i>	
Calvin Bao, Connor Baumler, Hal Daumé Iii and Marine Carpuat	25312
<i>UniSpeaker: A Unified Approach for Multimodality-driven Speaker Generation</i>	
Zhengyan Sheng, Zhihao Du, Heng Lu, ShiLiang Zhang and Zhen-Hua Ling	25331

<i>On the Fine-Grained Planning Abilities of VLM Web Agents</i> Surgan Jandial, Yinong Oliver Wang, Andrea Bajcsy and Fernando De la Torre	25347
<i>InterFeedback: Unveiling Interactive Intelligence of Large Multimodal Models with Human Feedback</i> Henry Hengyuan Zhao, Wenqi Pei, Yifei Tao, Haiyang Mei and Mike Zheng Shou	25381
<i>ReFLAIR: Enhancing Multimodal Reasoning via Structured Reflection and Reward-Guided Learning</i> Jiazhou Ji and Xinru Lu	25401
<i>ControlText: Unlocking Controllable Fonts in Multilingual Text Rendering without Font Annotations</i> Bowen Jiang, Yuan Yuan, Xinyi Bai, Zhuoqun Hao, Alyson Yin, Yaojie Hu, Wenyu Liao, Lyle Ungar and Camillo Jose Taylor	25414
<i>STA-CoT: Structured Target-Centric Agentic Chain-of-Thought for Consistent Multi-Image Geological Reasoning</i> Beibei Yu, Tao Shen and Ling Chen	25426
<i>Can Language Models Follow Multiple Turns of Entangled Instructions?</i> Chi Han, Xin Liu, Haodong Wang, Shiyang Li, Jingfeng Yang, Haoming Jiang, Zhengyang Wang, Qingyu Yin, Liang Qiu, Changlong Yu, Yifan Gao, Zheng Li, Bing Yin, Jingbo Shang and Heng Ji	25445
<i>How to Generalize the Detection of AI-Generated Text: Confounding Neurons</i> Claudio Borile and Carlo Abrate	25461
<i>SparsePO: Controlling Preference Alignment of LLMs via Sparse Token Masks</i> Fenia Christopoulou, Ronald Cardenas, Gerasimos Lampouras, Haitham Bou Ammar and Jun Wang	25477
<i>We Argue to Agree: Towards Personality-Driven Argumentation-Based Negotiation Dialogue Systems for Tourism</i> Priyanshu Priya, Saurav Dudhate, Desai Vishesh Yasheshbhai and Asif Ekbal	25504
<i>Towards the Roots of the Negation Problem: A Multilingual NLI Dataset and Model Scaling Analysis</i> Tereza Vrabcová, Marek Kadlčík, Petr Sojka, Michal Štefánik and Michal Spiegel	25537
<i>Improving the Language Understanding Capabilities of Large Language Models Using Reinforcement Learning</i> Sai Ashish Somayajula, Bokai Hu, Qi Cao, Xin Pan and Pengtao Xie	25552
<i>HATECAT-TR: A Hate Speech Span Detection and Categorization Dataset for Turkish</i> Hasan Kerem Şeker, Gökçe Uludoğan, Pelin Önal and Arzucan Özgür	25568
<i>DM-Codec: Distilling Multimodal Representations for Speech Tokenization</i> Md Mubtasim Ahasan, Md Fahim, Tasnim Mohiuddin, Akmmahbubur Rahman, Aman Chadha, Tariq Iqbal, M Ashraful Amin, Md Mofijul Islam and Amin Ahsan Ali	25580
<i>LCAN: A Label-Aware Contrastive Attention Network for Multi-Intent Recognition and Slot Filling in Task-Oriented Dialogue Systems</i> Shuli Zhang, Zhiqiang You, Xiao Xiang Qi, Peng Liu, Gaode Wu, Kan Xia and Shenguang Huang	25603
<i>Low-Resource Languages LLM Disinformation is Within Reach: The Case of Walliserdeutsch</i> Andrei Kucharavy, Sherine Seppey, Cyril Vallez, Dimitri Percia David and Ljiljana Dolamic	25613

<i>Exploring and Controlling Diversity in LLM-Agent Conversation</i> KuanChao Chu, Yi-Pei Chen and Hideki Nakayama	25626
<i>Agentic-ToM: Cognition-Inspired Agentic Processing For Enhancing Theory of Mind Reasoning</i> Sneheel Sarangi, Chetan Talele and Hanan Salam	25645
<i>Can We Edit LLMs for Long-Tail Biomedical Knowledge?</i> Xinhao Yi, Jake Lever, Kevin Bryson and Zaiqiao Meng	25662
<i>GeoPQA: Bridging the Visual Perception Gap in MLLMs for Geometric Reasoning</i> Guizhen Chen, Weiwen Xu, Hao Zhang, Hou Pong Chan, Deli Zhao, Anh Tuan Luu and Yu Rong 25680	
<i>CM-Align: Consistency-based Multilingual Alignment for Large Language Models</i> Xue Zhang, Yunlong Liang, Fandong Meng, Songming Zhang, Yufeng Chen, Jinan Xu and Jie Zhou	25689
<i>Cache Saver: A Modular Framework for Efficient, Affordable, and Reproducible LLM Inference</i> Nearchos Potamitis, Lars Henning Klein, Bardia Mohammadi, Chongyang Xu, Attreyee Mukher- jee, Niket Tandon, Laurent Bindschaedler and Akhil Arora	25703
<i>Evaluating Cultural Knowledge and Reasoning in LLMs Through Persian Allusions</i> Melika Nobakhtian, Yadollah Yaghoobzadeh and Mohammad Taher Pilehvar	25725
<i>Evolving Stances on Reproducibility: A Longitudinal Study of NLP and ML Researchers' Views and Experience of Reproducibility</i> Craig Thomson, Ehud Reiter, João Sedoc and Anya Belz	25738
<i>KAHAN: Knowledge-Augmented Hierarchical Analysis and Narration for Financial Data Narration</i> Yajing Yang, Tony Deng and Min-Yen Kan	25761