

**40th AAAI Conference on Artificial
Intelligence (AAAI-26), 38th
Conference on Innovative Applications
of Artificial Intelligence (IAAI-26), and
16th Symposium on Educational
Advances in Artificial Intelligence
(EAAI-26)**

Volume 37: AAAI Technical Tracks

- Natural Language Processing II

Singapore
20-27 January 2026

ISBN: 979-8-3313-3411-6

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© (2026) by Association for the Advancement of Artificial Intelligence
All rights reserved.

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

For permission requests, please contact Association for the Advancement of Artificial Intelligence
at the address below.

Association for the Advancement of Artificial Intelligence
2275 East Bayshore Road
Suite 160
Palo Alto, California 94303
USA

Phone: 1-650-328-3123
Fax: 1-650-321-4457

<https://aaai.org/Press/press.php>

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

AAAI TECHNICAL TRACK ON NATURAL LANGUAGE PROCESSING II

Mitigating Error Accumulation in Knowledge Editing for Multi-Hop Question Answering	30862
<i>Jiaxin Guo, Hao Sun, Wenhao Zhang, Xuanbo Fan, Yan Zhang</i>	
Seeing is Believing: Rich-Context Hallucination Detection for MLLMs Via Backward Visual Grounding.....	30871
<i>Pinxue Guo, Chongruo Wu, Xinyu Zhou, Lingyi Hong, Zhaoyu Chen, Jinglun Li, Kaixun Jiang, Sen-Ching Samson Cheung, Wei Zhang, Wenqiang Zhang</i>	
Extracting Events Like Code: A Multi-Agent Programming Framework for Zero-Shot Event Extraction	30880
<i>Quanjiang Guo, Sijie Wang, Jinchuan Zhang, Ben Zhang, Zhao Kang, Ling Tian, Ke Yan</i>	
MCP-AgentBench: Evaluating Real-World Language Agent Performance with MCP-Mediated Tools.....	30888
<i>Zikang Guo, Benfeng Xu, Chiwei Zhu, Wentao Hong, Xiaorui Wang, Zhendong Mao</i>	
MoSE: Hierarchical Self-Distillation Enhances Early Layer Embeddings.....	30897
<i>Andrea Gurioli, Federico Pennino, Joao Monteiro, Maurizio Gabbrielli</i>	
LaTeX2Layout: High-Fidelity, Scalable Document Layout Annotation Pipeline for Layout Detection	30907
<i>Feijiang Han, Zelong Wang, Bowen Wang, Xinxin Liu, Skyler Cheung, Delip Rao, Chris Callison-Burch, Lyle Ungar</i>	
Multi-Granularity Interactive Attention Framework for Residual Hierarchical Pronunciation Assessment.....	30916
<i>Hong Han, Hao-Chen Pei, Zhao-Zheng Nie, Xin Luo, Xin-Shun Xu</i>	
Uncertainty-Based Methods for Automated Process Reward Data Construction and Output Aggregation in Mathematical Reasoning	30925
<i>Jiuzhou Han, Wray Buntine, Ehsan Shareghi</i>	
Uncovering and Mitigating Transient Blindness in Multimodal Model Editing	30934
<i>Xiaoqi Han, Ru Li, Ran Yi, Hongye Tan, Zhuomin Liang, Victor Gutierrez Basulto, Jeff Z. Pan</i>	
Fact2Fiction: Targeted Poisoning Attack to Agentic Fact-Checking System	30943
<i>Haorui He, Yupeng Li, Bin Benjamin Zhu, Dacheng Wen, Reynold Cheng, Francis C. M. Lau</i>	
MentalGuide: Towards Multi-Turn, State-Aware and Strategy-Driven Conversations for Mental Health Support.....	30951
<i>Jinwei He, Feng Lu</i>	
PaSE: Prototype-Aligned Calibration and Shapley-Based Equilibrium for Multimodal Sentiment Analysis.....	30960
<i>Kang He, Boyu Chen, Yuzhe Ding, Fei Li, Chong Teng, Donghong Ji</i>	
ConInstruct: Evaluating Large Language Models on Conflict Detection and Resolution in Instructions	30969
<i>Xingwei He, Qianru Zhang, Pengfei Chen, Guanhua Chen, Linlin Yu, Yuan Yuan, Siu-Ming Yiu</i>	

TRACE: Trajectory-Based Activation Change Estimation for Task-Specific Data Selection	30978
<i>Ye He, Shangzhan Li, Yuxin Zhou, Qi Shi</i>	
FinMathBench: A Formula-Driven Benchmark for Evaluating LLMs' Math Reasoning Capabilities in Finance	30987
<i>Yi He, Ping Wang, Shiqiang Xiong, Chao Chen, Haixiang Hu</i>	
Exploring Depth Generalization in Large Language Models for Solving Recursive Logic Tasks	30996
<i>Zhiyuan He</i>	
REINA: Regularized Entropy Information-Based Loss for Efficient Simultaneous Speech Translation.....	31005
<i>Nameer Hirschkind, Joseph Liu, Xiao Yu, Mahesh Kumar Nandwana</i>	
Format Matters: The Robustness of Multimodal LLMs in Reviewing Evidence from Tables and Charts	31014
<i>Xanh Ho, Yun-Ang Wu, Sunisth Kumar, Florian Boudin, Atsuhiko Takasu, Akiko Aizawa</i>	
Benchmarking LLMs' Mathematical Reasoning with Unseen Random Variables Questions	31023
<i>Zijin Hong, Hao Wu, Su Dong, Junnan Dong, Yilin Xiao, Yujing Zhang, Zhu Wang, Feiran Huang, Linyi Li, Hongxia Yang, Xiao Huang</i>	
Cog-RAG: Cognitive-Inspired Dual-Hypergraph with Theme Alignment Retrieval-Augmented Generation	31032
<i>Hao Hu, Yifan Feng, Ruoxue Li, Rundong Xue, Xingliang Hou, Zhiqiang Tian, Yue Gao, Shaoyi Du</i>	
End-To-End Contrastive Language-Speech Pretraining Model for Long-Form Spoken Question Answering	31041
<i>Jiliang Hu, Zuchao Li, Baoyuan Qi, Guoming Liu, Ping Wang</i>	
AgentMental: An Interactive Multi-Agent Framework for Explainable and Adaptive Mental Health Assessment	31050
<i>Jinpeng Hu, Ao Wang, Qianqian Xie, Zhuo Li, Hui Ma, Dan Guo</i>	
DUP: Detection-Guided Unlearning for Backdoor Purification in Language Models	31059
<i>Man Hu, Yahui Ding, Yatao Yang, Liangyu Chen, Yanhao Jia, Shuai Zhao</i>	
TaxReasoning: Benchmarking Knowledge-Intensive Mathematical Reasoning with Evolving Tax Laws	31068
<i>Nan Hu, Yike Wu, Jiaye Li, Huikang Hu, Guilin Qi, Songlin Zhai, Yongrui Chen, Tianxing Wu, Tongtong Wu, Jiaoyan Chen, Jeff Z. Pan</i>	
On the Evaluation of Capability Estimation Methods for Large Language Models.....	31077
<i>Qiang Hu, Jin Wen, Yao Zhang, Maxime Cordy, Yongqiang Lyu</i>	
SAR: A Structure-Aligned Reasoning Framework for Temporal Knowledge Graph Question Answering	31086
<i>Qianyi Hu, Jiaxue Liu, Xinhui Tu, Shoujin Wang</i>	
Iterative Multi-Granular RAG with Contextual Hierarchical Graph.....	31095
<i>Yanli Hu, Teng Liu, Zhuangyi Zhou, Weixin Zeng, Zhen Tan, Xiang Zhao</i>	
Beyond Perplexity: Let the Reader Select Retrieval Summaries Via Spectrum Projection Score.....	31104
<i>Zhanghao Hu, Qinglin Zhu, Siya Qi, Yulan He, Hanqi Yan, Lin Gui</i>	

Model Editing as a Double-Edged Sword: Steering Agent Behavior Toward Beneficence Or Harm	31113
<i>Baixiang Huang, Zhen Tan, Haoran Wang, Zijie Liu, Dawei Li, Ali Payani, Huan Liu, Tianlong Chen, Kai Shu</i>	
HAPO: Training Language Models to Reason Concisely Via History-Aware Policy Optimization	31122
<i>Chengyu Huang, Zhengxin Zhang, Claire Cardie</i>	
RoSE: A Role Correlation Structure-Enhanced Model for Multi-Event Argument Extraction	31131
<i>Geting Huang, Jilong Zhang, Kai Zhou, Zhang Yi, Xiuyuan Xu</i>	
A Reasoning Paradigm for Named Entity Recognition.....	31140
<i>Hui Huang, Yanping Chen, Ruizhang Huang, Chuan Lin, Yongbin Qin</i>	
Long-Form RewardBench: Evaluating Reward Models for Long-Form Generation	31149
<i>Hui Huang, Yancheng He, Wei Liu, Muyun Yang, Jiaheng Liu, Kehai Chen, Bing Xu, Conghui Zhu, Hailong Cao, Tiejun Zhao</i>	
Think-J: Learning to Think for Generative LLM-As-A-Judge.....	31158
<i>Hui Huang, Yancheng He, Hongli Zhou, Rui Zhang, Wei Liu, Weixun Wang, Jiaheng Liu, Wenbo Su</i>	
EMODIS: A Benchmark for Context-Dependent Emoji Disambiguation in Large Language Models	31167
<i>Jiacheng Huang, Ning Yu, Xiaoyin Yi</i>	
Efficient Reasoning for Large Reasoning Language Models Via Certainty-Guided Reflection Suppression	31176
<i>Jiameng Huang, Baijiong Lin, Guhao Feng, Jierun Chen, Di He, Lu Hou</i>	
Whispering Agents: A Event-Driven Covert Communication Protocol for the Internet of Agents	31185
<i>Kaibo Huang, Yukun Wei, Wanshengwu , Tianhua Zhang, Zhongliang Yang, Linna Zhou</i>	
QiMeng-CRUX: Narrowing the Gap Between Natural Language and Verilog Via Core Refined Understanding eXpression.....	31193
<i>Lei Huang, Rui Zhang, Jiaming Guo, Yang Zhang, Di Huang, Shuyao Cheng, Pengwei Jin, Chongxiao Li, Zidong Du, Xing Hu, Yunji Chen, Qi Guo</i>	
Relink: Constructing Query-Driven Evidence Graph On-The-Fly for GraphRAG.....	31202
<i>Manzong Huang, Chenyang Bu, Yi He, Xingrui Zhuo, Xindong Wu</i>	
Hybrid Routing for a Mixture of LoRA Experts	31211
<i>Yitong Huang, Ziqi Yang, Zihui Wang, Jianzhong Qi, Rongshan Yu, Xiaoliang Fan, Cheng Wang</i>	
SPA: Achieving Consensus in LLM Alignment Via Self-Priority Optimization	31220
<i>Yue Huang, Xiangqi Wang, Xiangliang Zhang</i>	
Mem-PAL: Towards Memory-Based Personalized Dialogue Assistants for Long-Term User-Agent Interaction.....	31229
<i>Zhaopei Huang, Qifeng Dai, Guozheng Wu, Xiaopeng Wu, Xubin Li, Tiezheng Ge, Wenxuan Wang, Qin Jin</i>	
BiCycle: Group-Wise Recursive Transformer Based on ASR Mechanism	31238
<i>Min Ho Jang, Eun Seo Seo, Jin Young Kim, Hyeongsoo Lim, Ji Won Yoon</i>	
Do Language Models Associate Sound with Meaning? a Multimodal Study of Sound Symbolism.....	31247
<i>Jinhong Jeong, Sunghyun Lee, Jaeyoung Lee, Seonah Han, Youngjae Yu</i>	

Bootstrapping LLMs Via Preference-Based Policy Optimization	31256
<i>Chen Jia</i>	
AutoTool: Efficient Tool Selection for Large Language Model Agents	31265
<i>Jingyi Jia, Qinbin Li</i>	
LiteLong: Resource-Efficient Long-Context Data Synthesis for LLMs.....	31274
<i>Junlong Jia, Xing Wu, Chaochen Gao, Ziyang Chen, Zijia Lin, Zhongzhi Li, Weinong Wang, Haotian Xu, Donghui Jin, Debing Zhang, Binghui Guo</i>	
MEML-GRPO: Heterogeneous Multi-Expert Mutual Learning for RLVR Advancement	31283
<i>Weitao Jia, Jinghui Lu, Haiyang Yu, Siqi Wang, Guozhi Tang, An-Lan Wang, Weijie Yin, Dingkang Yang, Yuxiang Nie, Bin Shan, Hao Feng, Irene Li, Kun Yang, Han Wang, Jingqun Tang, Teng Fu, Changhong Jin, Chao Feng, Xiaohui Lv, Can Huang</i>	
MDF: A Modality-Aware Disentanglement and Fusion Framework for Multimodal Sentiment Analysis.....	31292
<i>Zhongquan Jian, Wenhan Lv, Yanhao Chen, Guanran Luo, Wentao Qiu, Shaopan Wang, Bingbing Hu, Qingqiang Wu</i>	
A Principle-Driven Adaptive Policy for Group Cognitive Stimulation Dialogue for Elderly with Cognitive Impairment.....	31301
<i>Jiyue Jiang, Yanyu Chen, Pengan Chen, Kai Liu, Jingqi Zhou, Zheyong Zhu, He Hu, Fei Ma, Qi Tian, Chuan Wu</i>	
Consensus-Aligned Neuron Efficient Fine-Tuning Large Language Models for Multi-Domain Machine Translation	31310
<i>Shuting Jiang, Ran Song, Yuxin Huang, Yan Xiang, Yantuan Xian, Shengxiang Gao, Zhengtao Yu</i>	
MedS ³ : Towards Medical Slow Thinking with Self-Evolved Soft Dual-Sided Process Supervision	31319
<i>Shuyang Jiang, Yusheng Liao, Zhe Chen, Ya Zhang, Yanfeng Wang, Yu Wang</i>	
Importance-Aware Data Selection for Efficient LLM Instruction Tuning.....	31328
<i>Tingyu Jiang, Shen Li, Yiyao Song, Lan Zhang, Hualei Zhu, Yuan Zhao, Xiaohang Xu, Kenjiro Taura, Hao Henry Wang</i>	
From Chaos to Clarity: A Knowledge Graph-Driven Audit Dataset Generation Framework for LLM Unlearning.....	31337
<i>Weipeng Jiang, Juan Zhai, Shiqing Ma, Ziyang Lei, Xiaofei Xie, Yige Wang, Chao Shen</i>	
Large Language Model Unlearning for Source Code	31346
<i>Xue Jiang, Yihong Dong, Huangzhao Zhang, Tangxinyu Wang, Zheng Fang, Yingwei Ma, Rongyu Cao, Binhua Li, Zhi Jin, Wenpin Jiao, Yongbin Li, Ge Li</i>	
EduGuardBench: A Holistic Benchmark for Evaluating the Pedagogical Fidelity and Adversarial Safety of LLMs as Simulated Teachers	31356
<i>Yilin Jiang, Mingzi Zhang, Xuanyu Yin, Sheng Jin, Suyu Lu, Zuocan Ying, Zengyi Yu, Xiangjie Kong</i>	
Difficulty is Not Enough: Curriculum Learning for LLMs Fine-Tuning Must Consider Utility.....	31365
<i>Zishang Jiang, Jinyi Han, Tingyun Li, Xinyi Wang, Sihang Jiang, Xiaojun Meng, Jiansheng Wei, Jiaqing Liang, Yanghua Xiao</i>	
Embracing Positional Bias in Multiple-Choice Question Answering Via Permutation Equivariant Neural Networks.....	31374
<i>Chengyu Jiao, Siyin Huang, Yu Zhang</i>	

Causal-ERC: A Multimodal Framework with Causal Prompting for Emotion Recognition in Conversations with Large Language Models	31383
<i>Ran Jing, Geng Tu, Yice Zhang, Ruifeng Xu</i>	
Model Whisper: Steering Vectors Unlock Large Language Models' Potential in Test-Time	31392
<i>Xinyue Kang, Diwei Shi, Li Chen</i>	
WikiMAG: A Multi-Agent Guided Framework for Generating Structured Wikipedia-Like Articles	31401
<i>Xiuli Kang, Yinlong Xiao, Minghao Hu, Yuan Huang, Bin Mao, Ming Wang, Fang Wang, Zhunchen Luo, Wei Luo, Guotong Geng</i>	
SampurNER: Fine-Grained Named Entity Recognition Dataset for 22 Indian Languages	31410
<i>Prachuryya Kaushik, Ashish Anand</i>	
CATCH: A Controllable Theme Detection Framework with Contextualized Clustering and Hierarchical Generation.....	31419
<i>Rui Ke, Jiahui Xu, Shenghao Yang, Kuang Wang, Feng Jiang, Haizhou Li</i>	
Test-Time Diverse Reasoning by Riemannian Activation Steering.....	31429
<i>Ly Tran Ho Khanh, Dongxuan Zhu, Man-Chung Yue, Viet Anh Nguyen</i>	
Rethinking the Sampling Criteria in Reinforcement Learning for LLM Reasoning: A Competence-Difficulty Alignment Perspective	31438
<i>Deyang Kong, Qi Guo, Xiangyu Xi, Wei Wang, Jingang Wang, Xunliang Cai, Shikun Zhang, Wei Ye</i>	
Backdooring Rationalization	31447
<i>Lingxiao Kong, Jiahui Jiang, Wenchao Xu, Lei Wu</i>	
LC3: Long Cross-Language Code Clone Detection Enhanced by Opcode Sequences and Affinity Aggregation.....	31456
<i>Xilin Lan, Huan Zhang, Yang Yang, Chengwu Xue, Li Kuang</i>	
Template-Theorems Graph Construction to Enhance Mathematical Reasoning Capabilities of LLM.....	31465
<i>Yarong Lan, Yajing Xu, Huajun Chen</i>	
Reinforce Trustworthiness in Multimodal Emotional Support System	31474
<i>Huy M. Le, Dat Tien Nguyen, Ngan T. T. Vo, Tuan D. Q. Nguyen, Nguyen Le Binh, Duy Minh Ho Nguyen, Daniel Sonntag, Lizi Liao, Binh T. Nguyen</i>	
Confidence-Guided Stepwise Model Routing for Cost-Efficient Reasoning	31483
<i>Sangmook Lee, Dohyung Kim, Hyukhun Koh, Nakyeong Yang, Kyomin Jung</i>	
The Curious Case of Analogies: Investigating Analogical Reasoning in Large Language Models.....	31492
<i>Taewho Lee, Minju Song, Chanwoong Yoon, Jungwoo Park, Jaewoo Kang</i>	
TaREx: Reinforcement Learning for Code-Driven Table Reasoning	31501
<i>Fangyu Lei, Jinxiang Meng, Yiming Huang, Shizhu He, Jun Zhao, Kang Liu</i>	
DegVoC: Revisiting Neural Vocoder from a Degradation Perspective.....	31510
<i>Andong Li, Tong Lei, Lingling Dai, Kai Li, Rilin Chen, Meng Yu, Xiaodong Li, Dong Yu, Chengshi Zheng</i>	
Language Drift in Multilingual Retrieval-Augmented Generation: Characterization and Decoding-Time Mitigation.....	31519
<i>Bo Li, Zhenghua Xu, Rui Xie</i>	

Modeling Uncertainty Trends for Timely Retrieval in Dynamic RAG	31527
<i>Bo Li, Tian Tian, Zhenghua Xu, Hao Cheng, Shikun Zhang, Wei Ye</i>	
HPSU: A Benchmark for Human-Level Perception in Real-World Spoken Speech Understanding	31536
<i>Chen Li, Peiji Yang, Yicheng Zhong, Jianxing Yu, Zhisheng Wang, Zihao Gou, Wenqing Chen, Jian Yin</i>	
Res-Bench: Benchmarking the Robustness of Multimodal Large Language Models to Dynamic Resolution Input	31545
<i>Chenxu Li, Zhicai Wang, Yuan Sheng, Xingyu Zhu, Yanbin Hao, Xiang Wang</i>	
CliCARE: Grounding Large Language Models in Clinical Guidelines for Decision Support Over Longitudinal Cancer Electronic Health Records	31554
<i>Dongchen Li, Jitao Liang, Wei Li, Xiaoyu Wang, Longbing Cao, Kun Yu</i>	
KVmix: Gradient-Based Layer Importance-Aware Mixed-Precision Quantization for KV Cache	31563
<i>Fei Li, Song Liu, Weiguo Wu, Shiqiang Nie, Jinyu Wang</i>	
K-STaR: Knowledge-Aware Self-Taught Reasoner.....	31573
<i>Guozheng Li, Xinyu Zhang</i>	
From Sampling to Cognition: Modeling Internal Cognitive Confidence in Language Models for Robust Uncertainty Calibration	31582
<i>Hao Li, Tao He, Jiafeng Liang, Zheng Chu, Ming Liu</i>	
Semi-Supervised Synthetic Data Generation with Fine-Grained Relevance Control for Short Video Search Relevance Modeling	31591
<i>Haoran Li, Zhiming Su, Junyan Yao, Enwei Zhang, Yang Ji, Yan Chen, Kan Zhou, Chao Feng, Jiao Ran</i>	
From Macro to Micro: Probing Dataset Diversity in Language Model Fine-Tuning	31600
<i>Haoyu Li, Xuhong Li, Yiming Dong, Kun Liu</i>	
From Single to Societal: Analyzing Persona-Induced Bias in Multi-Agent Interactions	31609
<i>Jiayi Li, Xiao Liu, Yansong Feng</i>	
LifeAlign: Lifelong Alignment for Large Language Models with Memory-Augmented Focalized Preference Optimization	31618
<i>Junsong Li, Jie Zhou, Bihao Zhan, Yutao Yang, Qianjun Pan, Shilian Chen, Tianyu Huai, Xin Li, Qin Chen, Liang He</i>	
WenetSpeech-Yue: A Large-Scale Cantonese Speech Corpus with Multi-Dimensional Annotation.....	31627
<i>Longhao Li, Zhao Guo, Hongjie Chen, Yuhang Dai, Ziyu Zhang, Hongfei Xue, Tianlun Zuo, Chengyou Wang, Shuiyuan Wang, Xin Xu, Hui Bu, Jie Li, Jian Kang, Binbin Zhang, Ruibin Yuan, Ziya Zhou, Wei Xue, Lei Xie</i>	
RecToM: A Benchmark for Evaluating Machine Theory of Mind in LLM-Based Conversational Recommender Systems	31636
<i>Mengfan Li, Xuanhua Shi, Yang Deng</i>	
Causal Tracing of Object Representations in Large Vision Language Models: Mechanistic Interpretability and Hallucination Mitigation	31645
<i>Qiming Li, Zekai Ye, Xiaocheng Feng, Weihong Zhong, Weitao Ma, Xiachong Feng</i>	

RAGFort: Dual-Path Defense Against Proprietary Knowledge Base Extraction in Retrieval-Augmented Generation.....	31654
<i>Qinfeng Li, Miao Pan, Ke Xiong, Ge Su, Zhiqiang Shen, Yan Liu, Sun Bing, Hao Peng, Xuhong Zhang</i>	
Do Not Merge My Model! Safeguarding Open-Source LLMs Against Unauthorized Model Merging	31662
<i>Qinfeng Li, Miao Pan, Jintao Chen, Fu Teng, Zhiqiang Shen, Ge Su, Hao Peng, Xuhong Zhang</i>	
From Hypothesis to Premises: LLM-Based Backward Logical Reasoning with Selective Symbolic Translation.....	31671
<i>Qingchuan Li, Mingyue Cheng, Zirui Liu, Daoyu Wang, Yuting Zeng, Tongxuan Liu</i>	
AdaFuse: Accelerating Dynamic Adapter Inference Via Token-Level Pre-Gating and Fused Kernel Optimization.....	31680
<i>Qiyang Li, Rui Kong, Yuchen Li, Hengyi Cai, Shuaiqiang Wang, Linghe Kong, Guihai Chen, Dawei Yin</i>	
Learning Beyond Domains: Misleading Prompts and Pseudo-Label Contrast for Text Domain Generalization	31689
<i>Qizhi Li, Xuyang Wang, Yingke Chen, Ming Yan, Dezhong Peng, Xi Peng, Xu Wang</i>	
OSVBench: Benchmarking LLMs on Specification Generation Tasks for Operating System Verification	31698
<i>Shangyu Li, Juyong Jiang, Tiancheng Zhao, Jiasi Shen</i>	
CoFact: Dynamic Coordination of Attention Heads for Improving Factual Consistency in LLMs	31708
<i>Shike Li, Xiaokai Wang, Xiaofeng Liu, Xin Tong, Hu Zhang</i>	
Automatic Paper Reviewing with Heterogeneous Graph Reasoning Over LLM-Simulated Reviewer-Author Debates	31717
<i>Shuaimin Li, Liyang Fan, Yufang Lin, Zeyang Li, Xian Wei, Shiwen Ni, Hamid Alinejad-Rokny, Min Yang</i>	
Jupiter: Enhancing LLM Data Analysis Capabilities Via Notebook and Inference-Time Value-Guided Search	31726
<i>Shuo Cheng Li, Yihao Liu, Silin Du, Wenxuan Zeng, Zhe Xu, Mengyu Zhou, Yeye He, Haoyu Dong, Shi Han, Dongmei Zhang</i>	
Step-GRPO: Enhancing Reasoning Quality and Efficiency Via Structured PRM-Based Reinforcement Learning	31734
<i>Weijie Li, Jin Wang, Liang-Chih Yu, Xuejie Zhang</i>	
ENCORE: Entropy-Guided Reward Composition for Multi-Head Safety Reward Models.....	31743
<i>Xiaomin Li, Xupeng Chen, Jingxuan Fan, Eric Hanchen Jiang, Mingye Gao</i>	
Semantic Volume: Quantifying and Detecting Both External and Internal Uncertainty in LLMs	31751
<i>Xiaomin Li, Zhou Yu, Ziji Zhang, Yingying Zhuang, Swair Shah, Narayanan Sadagopan, Anurag Beniwal</i>	
Selection of LLM Fine-Tuning Data Based on Orthogonal Rules	31760
<i>Xiaomin Li, Mingye Gao, Zhiwei Zhang, Chang Yue, Hong Hu</i>	

Author Index