

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 41: AAAI Technical Tracks

- Natural Language Processing VI

Singapore
20-27 January 2026

ISBN: 979-8-3313-3415-4

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 VI

WhisperDiari: A Whisper-Based Speaker Diarization Framework in Token Space Leveraging Semantic and Speaker Information for Better Text Adaptability.....	34477
<i>Yongkang Yin, Yuexian Zou</i>	
SR-KI: Scalable and Real-Time Knowledge Integration into LLMs Via Supervised Attention.....	34486
<i>Bohan Yu, Wei Huang, Kang Liu</i>	
ToC: Tree-Of-Claims Search with Multi-Agent Language Models.....	34495
<i>Shuyang Yu, Jianan Liang, Hui Hu</i>	
Discovering Decoupled Functional Modules in Large Language Models.....	34503
<i>Yanke Yu, Jin Li, Ying Sun, Ping Li, Zhefeng Wang, Yi Zheng</i>	
Language Model Distillation: A Temporal Difference Imitation Learning Perspective	34512
<i>Zishun Yu, Shangzhe Li, Xinhua Zhang</i>	
MCTS-SQL: Light-Weight LLMs Can Master the Text-To-SQL Through Monte Carlo Tree Search	34521
<i>Shuozi Yuan, Liming Chen, Miaomiao Yuan, Zhao Jin</i>	
Promoting Efficient Reasoning with Verifiable Stepwise Reward	34530
<i>Chuhuai Yue, Chengqi Dong, Yinan Gao, Hang He, Jiajun Chai, Wei Lin, Guojun Yin</i>	
MARS: Multimodal Adaptive Reasoning Model for Avoiding Overthinking	34539
<i>Tan Yue, Qiong Wu, Dongyan Zhao</i>	
Reducing the Scope of Language Models	34548
<i>David Yunis, Siyu Huo, Chulaka Gunasekara, Danish Contractor</i>	
ExPairT-LLM: Exact Learning for LLM Code Selection by Pairwise Queries.....	34557
<i>Tom Yuviler, Dana Drachler-Cohen</i>	
CAMA: Enhancing Mathematical Reasoning in Large Language Models with Causal Knowledge	34566
<i>Lei Zan, Keli Zhang, Ruichu Cai, Lujia Pan</i>	
Expert-Inspired Multi-Agent Coordination for Multi-Objective Molecular Optimization.....	34575
<i>Daojian Zeng, Tianle Li, Jiahao Yang, Jiakai Yi, Xieping Gao, Lincheng Jiang, Tengfei Ma, Xiangxiang Zeng</i>	
SlideTailor: Personalized Presentation Slide Generation for Scientific Papers	34584
<i>Wenzheng Zeng, Mingyu Ouyang, Langyuan Cui, Hwee Tou Ng</i>	
ToolACE-R: Model-Aware Iterative Training and Adaptive Refinement for Tool Learning	34593
<i>Xingshan Zeng, Weiwen Liu, Xu Huang, Zezhong Wang, Lingzhi Wang, Liangyou Li, Yasheng Wang, Lifeng Shang, Xin Jiang, Ruiming Tang, Qun Liu</i>	
MathSmith: Towards Extremely Hard Mathematical Reasoning by Forging Synthetic Problems with a Reinforced Policy	34602
<i>Shaoyang Zhan, Yanlin Lai, Ziyu Lu, Dahua Lin, Ziqing Yang, Fei Tan</i>	

JELV: A Judge of Edit-Level Validity for Evaluation and Automated Reference Expansion in Grammatical Error Correction.....	34611
<i>Yuhao Zhan, Yuqing Zhang, Jing Yuan, Qixiang Ma, Zhiqi Yang, Yu Gu, Zemin Liu, Fei Wu</i>	
Global-Local Confidence Fusion for Hallucination Detection in Mathematical Reasoning Task.....	34620
<i>Bo Zhang, Cong Gao, Linkang Yang, Bingxu Han, Minghao Hu, Zhunchen Luo, Guotong Geng, Xiaoying Bai, Jun Zhang, Wen Yao, Zhong Wang</i>	
Predicting Emergent Tool Use in LLMs Before it Emerges: A Proxy Perspective.....	34629
<i>Bo-Wen Zhang, Yan Yan, Guang Liu, Xu-Cheng Yin</i>	
Induce, Align, Predict: Zero-Shot Stance Detection Via Cognitive Inductive Reasoning.....	34638
<i>Bowen Zhang, Jun Ma, Fuqiang Niu, Li Dong, Jinzhou Cao, Genan Dai</i>	
RFKG-CoT: Relation-Driven Adaptive Hop-Count Selection and Few-Shot Path Guidance for Knowledge-Aware QA	34647
<i>Chao Zhang, Minghan Li, Tianrui Lv, Guodong Zhou</i>	
EA-VAE: Learning to Reconstruct Dysarthric Speech Via Variational Autoencoder with Encoding Alignment.....	34656
<i>Daipeng Zhang, Wenhuan Lu, Xianghu Yue, Hongcheng Zhang, Jianguo Wei</i>	
Learning to Generate and Extract: A Multi-Agent Collaboration Framework for Zero-Shot Document-Level Event Arguments Extraction.....	34665
<i>Guangjun Zhang, Hu Zhang, Yazhou Han, Yue Fan, Yuhang Shao, Hongye Tan, Ru Li</i>	
WikiREVIEW: A Multi-Perspective Review Framework for Automatic Wiki-Style Article Generation	34674
<i>Guo-Biao Zhang, Zhijing Wu, Tian Lan, Ding-Yuan Liu, Yu-Shi Zhu, Xian-Ling Mao</i>	
MoLoRA: Boosting LLM-Based End-To-End Speech Translation with Mixture of Low-Rank Experts.....	34683
<i>Hao Zhang, Yaqi Chen, Nianwen Si, Xukui Yang, Wenlin Zhang, Dan Qu</i>	
LLaVA-MS-PIT: Multi-Modal Schema-Guided Progressive Instruction Tuning for Multi-Modal Event Extraction	34692
<i>Hui Zhang, Po Hu, Wei Emma Zhang</i>	
AP2O-Coder: Adaptively Progressive Preference Optimization for Reducing Compilation and Runtime Errors in LLM-Generated Code.....	34701
<i>Jianqing Zhang, Wei Xia, Hande Dong, Qiang Lin, Jian Cao</i>	
Prune4Web: DOM Tree Pruning Programming for Web Agent	34710
<i>Jiayuan Zhang, Kaiquan Chen, Zhihao Lu, Enshen Zhou, Qian Yu, Jing Zhang</i>	
Steering Representations, Safeguarding Privacy: A Cross-Modal Privacy Protection Method for Generative AI	34719
<i>Jie Zhang, Chenxu Niu, Zhefeng Nan, Yangyan Xu, Jinta Weng</i>	
TAdaRAG: Task Adaptive Retrieval-Augmented Generation Via On-The-Fly Knowledge Graph Construction	34728
<i>Jie Zhang, Bo Tang, Wanzi Shao, Wenqiang Wei, Jihao Zhao, Jianqing Zhu, Zhiyu Li, Wen Xi, Zehao Lin, Feiyu Xiong, Yanchao Tan</i>	
Multi-Metric Preference Alignment for Generative Speech Restoration.....	34737
<i>Junan Zhang, Xueyao Zhang, Jing Yang, Yuancheng Wang, Fan Fan, Zhizheng Wu</i>	

LLM-CAS: Dynamic Neuron Perturbation for Real-Time Hallucination Correction	34746
<i>Jusheng Zhang, Ningyuan Liu, Yijia Fan, Zihao Huang, Qinglin Zeng, Kaitong Cai, Jian Wang, Keze Wang</i>	
MetaGDPO: Alleviating Catastrophic Forgetting with Metacognitive Knowledge Through Group Direct Preference Optimization	34755
<i>Lanxue Zhang, Yuqiang Xie, Fang Fang, Fanglong Dong, Rui Liu, Yanan Cao</i>	
CATS: Category-Aware Token-Level Steering for Training-Free Redundancy Reduction in Large Reasoning Models	34764
<i>Mengfei Zhang, Zhenglin Wang</i>	
Collaborative LLM Numerical Reasoning with Local Data Protection	34772
<i>Min Zhang, Yuzhe Lu, Yun Zhou, Panpan Xu, Lin Lee Cheong, Chang-Tien Lu, Haozhu Wang</i>	
ASKD: Reinforcement Learning-Style Knowledge Distillation with Quality-Adaptive Skewness	34781
<i>Mingjie Zhang, Xiaoling Zhou, Yuxiao Luo, Yiyu Liu, Shikun Zhang, Wei Ye</i>	
AI-Salesman: Towards Reliable Large Language Model Driven Telemarketing	34790
<i>Qingyu Zhang, Chunlei Xin, Xuanang Chen, Yaojie Lu, Hongyu Lin, Xianpei Han, Le Sun, Qing Ye, Qianlong Xie, Xingxing Wang</i>	
How Does Alignment Enhance LLMs' Multilingual Capabilities? a Language Neurons Perspective	34799
<i>Shimao Zhang, Zhejian Lai, Xiang Liu, Shuaijie She, Xiao Liu, Yeyun Gong, Shujian Huang, Jiajun Chen</i>	
Interpretable Reward Model Via Sparse Autoencoder.....	34808
<i>Shuyi Zhang, Wei Shi, Sihang Li, Jiayi Liao, Tao Liang, Hengxing Cai, Xiang Wang</i>	
L2-LoRA: Improving Low-Rank Adaptation with Layer-Specific Regularization	34817
<i>Xiang Zhang, Rui Xie, Shikun Zhang</i>	
Safety Alignment of Large Language Models Via Contrasting Safe and Harmful Distributions	34827
<i>Xiaoyun Zhang, Zhengyue Zhao, Wenxuan Shi, Kaidi Xu, Di Huang, Xing Hu</i>	
Token-Context Attention for NLI: An Alternative to Self-Attention	34836
<i>Xin Zhang, Victor S. Sheng</i>	
ShadeEdit: A Utility-Preserving and Defense-Evasive Knowledge Manipulation Attack in Federated LLMs	34845
<i>Xu Zhang, Hangcheng Liu, Shangwei Guo, Shudong Zhang, Tianwei Zhang, Tao Xiang</i>	
UDA: Unsupervised Debiasing Alignment for Pair-Wise LLM-As-A-Judge	34854
<i>Yang Zhang, Cunxiang Wang, Lindong Wu, Wenbo Yu, Yidong Wang, Guangsheng Bao, Jie Tang</i>	
LeanRAG: Knowledge-Graph-Based Generation with Semantic Aggregation and Hierarchical Retrieval	34862
<i>Yaoze Zhang, Rong Wu, Pinlong Cai, Xiaoman Wang, Guohang Yan, Song Mao, Ding Wang, Botian Shi</i>	
The Avengers: A Routing Recipe for Collective Intelligence in Language Models	34870
<i>Yiqun Zhang, Hao Li, Chenxu Wang, Linyao Chen, Qiaosheng Zhang, Peng Ye, Shi Feng, Xinrun Wang, Jia Xu, Lei Bai, Shuyue Hu</i>	
MSME: A Multi-Stage Multi-Expert Framework for Zero-Shot Stance Detection.....	34879
<i>Yuanshuo Zhang, Aohua Li, Bo Chen, Jingbo Sun, Xiaobing Zhao</i>	

Syllogism-Inspired TableQA: Evidentialization Makes Decomposition Reasoning and Answer Verification More Reliable	34888
<i>Zhe Zhang, Lili Bai, Chaopeng Guo, Jie Song</i>	
SCOPE: Intrinsic Semantic Space Control for Mitigating Copyright Infringement in LLMs.....	34897
<i>Zhenliang Zhang, Xinyu Hu, Xiaojun Wan</i>	
When Instinct Guides and Insight Grounds: Staged RL Training for LLM Agents.....	34906
<i>Zijing Zhang, Boning Zhang</i>	
TASE: Token Awareness and Structured Evaluation for Multilingual Language Models	34915
<i>Chenzhuo Zhao, Xinda Wang, Yue Huang, Junting Lu, Ziqian Liu</i>	
Reflect Then Learn: Active Prompting for Information Extraction Guided by Introspective Confusion	34923
<i>Dong Zhao, Yadong Wang, Xiang Chen, Chenxi Wang, Hongliang Dai, Chuanxing Geng, Shengzhong Zhang, Shao-Yuan Li, Sheng-Jun Huang</i>	
GenPRM: Scaling Test-Time Compute of Process Reward Models Via Generative Reasoning	34932
<i>Jian Zhao, Runze Liu, Kaiyan Zhang, Zhimu Zhou, Junqi Gao, Dong Li, Jiafei Lyu, Zhouyi Qian, Biqing Qi, Xiu Li, Bowen Zhou</i>	
Beyond Step Pruning: Information Theory Based Step-Level Optimization for Self-Refining Large Language Models	34941
<i>Jinman Zhao, Erxue Min, Hui Wu, Ziheng Li, Zexu Sun, Hengyi Cai, Shuaiqiang Wang, Xu Chen, Gerald Penn</i>	
SABER: Switchable and Balanced Training for Efficient LLM Reasoning.....	34950
<i>Kai Zhao, Yanjun Zhao, Jiaming Song, Shien He, Lusheng Zhang, Qiang Zhang, Tianjiao Li</i>	
Making Every Head Count: Sparse Attention Without the Speed-Performance Trade-Off.....	34959
<i>Mingkuan Zhao, Wentao Hu, Jiayin Wang, Xin Lai, Tianchen Huang, Yuheng Min, Rui Yan, Xiaoyan Zhu</i>	
MMRAG-RFT: Two-Stage Reinforcement Fine-Tuning for Explainable Multi-Modal Retrieval-Augmented Generation.....	34968
<i>Shengwei Zhao, Jingwen Yao, Sitong Wei, Linhai Xu, Yuying Liu, Dong Zhang, Zhiqiang Tian, Shaoyi Du</i>	
Trade-Offs in Large Reasoning Models: An Empirical Analysis of Deliberative and Adaptive Reasoning Over Foundational Capabilities	34976
<i>Weixiang Zhao, Xingyu Sui, Jiahe Guo, Yulin Hu, Yang Deng, Yanyan Zhao, Xuda Zhi, Yongbo Huang, Hao He, Wanxiang Che, Ting Liu, Bing Qin</i>	
AgentCDM: Enhancing Multi-Agent Collaborative Decision-Making Via ACH-Inspired Structured Reasoning	34985
<i>Xuyang Zhao, Shiwan Zhao, Hualong Yu, Liting Zhang, Qicheng Li</i>	
LaF-GRPO: In-Situ Navigation Instruction Generation for the Visually Impaired Via GRPO with LLM-As-Follower Reward.....	34994
<i>Yi Zhao, Siqi Wang, Jing Li</i>	
Don't Start Over: A Cost-Effective Framework for Migrating Personalized Prompts Between LLMs.....	35003
<i>Ziyi Zhao, Chongming Gao, Yang Zhang, Haoyan Liu, Weinan Gan, Hui Feng Guo, Yong Liu, Fuli Feng</i>	

Rethinking the Reliability of Multi-Agent System: A Perspective from Byzantine Fault Tolerance.....	35012
<i>Lifan Zheng, Jiawei Chen, Qinghong Yin, Jingyuan Zhang, Xinyi Zeng, Yu Tian</i>	
Say More with Less: Variable-Frame-Rate Speech Tokenization Via Adaptive Clustering and Implicit Duration Coding	35021
<i>Rui-Chen Zheng, Wenrui Liu, Hui-Peng Du, Qinglin Zhang, Chong Deng, Qian Chen, Wen Wang, Yang Ai, Zhen-Hua Ling</i>	
M3UCD: A Multi-Task Multimodal Metaphor Understanding Challenge Dataset for LLMs.....	35030
<i>Tianlong Zheng, Yating Yang, Rui Dong, Bo Ma, Lei Wang, Xi Zhou, Siru Miao, Turghun Osman</i>	
Navigating Through Paper Flood: Advancing LLM-Based Paper Evaluation Through Domain- Aware Retrieval and Latent Reasoning	35041
<i>Wuqiang Zheng, Yiyang Xu, Xinyu Lin, Chongming Gao, Wenjie Wang, Fuli Feng</i>	
Graph-Augmented and Over-Smoothing-Resistant Contrastive Clustering for Short Text.....	35050
<i>Zijian Zheng, Tao Ai, Yonghe Lu</i>	
Compress-Then-Rank: Faster and Better Listwise Reranking with Large Language Models Via Ranking-Aware Passage Compression	35059
<i>Zhewei Zhi, Yingyi Zhang, Yizhen Jing, Xianneng Li, Jianing Liu, Huajie Liu, Yongliang Ding</i>	
Learning from Guidelines: Structured Prompt Optimization for Expert Annotation Tasks.....	35068
<i>Wenliang Zhong, Haiqing Li, Thao M. Dang, Feng Jiang, Hehuan Ma, Yuzhi Guo, Jean Gao, Junzhou Huang</i>	
P2S: Probabilistic Process Supervision for General-Domain Reasoning Question Answering.....	35076
<i>Wenlin Zhong, Chengyuan Liu, Yiquan Wu, Bovin Tan, Changlong Sun, Yi Wang, Xiaozhong Liu, Kun Kuang</i>	
Lost in Benchmarks? Rethinking Large Language Model Benchmarking with Item Response Theory	35085
<i>Hongli Zhou, Hui Huang, Ziqing Zhao, Lvyuan Han, Huicheng Wang, Kehai Chen, Muyun Yang, Wei Bao, Jian Dong, Bing Xu, Conghui Zhu, Hailong Cao, Tiejun Zhao</i>	
HyperGLLM: An Efficient Framework for Endpoint Threat Detection Via Hypergraph-Enhanced Large Language Models	35094
<i>Hongyi Zhou, Jianfeng Pan, Min Peng, Shaomang Huang, Hanzhong Zheng</i>	
LLM Collaborative Filtering: User-Item Graph as New Language.....	35103
<i>Huachi Zhou, Yujing Zhang, Hao Chen, Qinggang Zhang, Qijie Shen, Feiran Huang, Xiao Huang</i>	
DIFFA: Large Language Diffusion Models Can Listen and Understand	35112
<i>Jiaming Zhou, Hongjie Chen, Shiwang Zhao, Jian Kang, Jie Li, Enzhi Wang, Yujie Guo, Haoqin Sun, Hui Wang, Aobo Kong, Yong Qin, Xuelong Li</i>	
STaR: Sensitive Trajectory Regulation for Unlearning in Large Reasoning Models	35121
<i>Jingjing Zhou, Gaoxiang Cong, Li Su, Liang Li</i>	
What to Ask Next? Probing the Imaginative Reasoning of LLMs with TurtleSoup Puzzles	35130
<i>Mengtao Zhou, Sifan Wu, Huan Zhang, Qi Sima, Bang Liu</i>	
IndexTTS2: A Breakthrough in Emotionally Expressive and Duration-Controlled Auto-Regressive Zero-Shot Text-To-Speech.....	35139
<i>Siyi Zhou, Yiquan Zhou, Yi He, Xun Zhou, Jinchao Wang, Wei Deng, Jingchen Shu</i>	

Note2Chat: Improving LLMs for Multi-Turn Clinical History Taking Using Medical Notes	35149
<i>Yang Zhou, Zhenting Sheng, Mingrui Tan, Yuting Song, Jun Zhou, Yu Heng Kwan, Lian Leng Low, Yang Bai, Yong Liu</i>	
Do Retrieval Augmented Language Models Know When They Don't Know?.....	35158
<i>Youchao Zhou, Heyan Huang, Yicheng Liu, Rui Dai, Xinglin Wang, Xingchen Zhang, Shumin Shi, Yang Deng</i>	
AncientBench: Towards Comprehensive Evaluation on Excavated and Transmitted Chinese Corpora.....	35167
<i>Zhihan Zhou, Daqian Shi, Rui Song, Lida Shi, Xiaolei Diao, Hao Xu</i>	
ResMAS: Resilience Optimization in LLM-Based Multi-Agent Systems	35176
<i>Zhilun Zhou, Zihan Liu, Jiahe Liu, Qingyu Shao, Yihan Wang, Kun Shao, Depeng Jin, Fengli Xu</i>	
Evaluating, Synthesizing, and Enhancing for Customer Support Conversation.....	35185
<i>Jie Zhu, Huaixia Dou, Junhui Li, Lifan Guo, Feng Chen, Chi Zhang, Fang Kong</i>	
In-Token Rationality Optimization: Towards Accurate and Concise LLM Reasoning Via Self-Feedback.....	35195
<i>Mingye Zhu, Yi Liu, Zheren Fu, Quan Wang, Yongdong Zhang</i>	
SubGCatch: Accelerating Graph-Based RAG with Subgraph-Level KV Cache	35204
<i>Qiuyu Zhu, Liang Zhang, Qianxiong Xu, Cheng Long, Jie Zhang</i>	
Outlier Matters: Efficient Long-To-Short Reasoning Via Outlier-Guided Model Merging.....	35213
<i>Qiyuan Zhu, Dezhi Li, Lujun Li, Xiaoyu Qin, Wei Li, Hao Gu, Hua Xu, Sirui Han, Yike Guo</i>	
Beyond Chains: Bridging Large Language Models and Knowledge Bases in Complex Question Answering	35222
<i>Yihua Zhu, Qianying Liu, Akiko Aizawa, Hidetoshi Shimodaira</i>	
REAP: Enhancing RAG with Recursive Evaluation and Adaptive Planning for Multi-Hop Question Answering	35230
<i>Yijie Zhu, Haojie Zhou, Wanting Hong, Tailin Liu, Ning Wang</i>	
Why Do Open-Source LLMs Struggle with Data Analysis? a Systematic Empirical Study	35239
<i>Yuqi Zhu, Yi Zhong, Jintian Zhang, Ziheng Zhang, Shuofei Qiao, Yujie Luo, Lun Du, Da Zheng, Ningyu Zhang, Huajun Chen</i>	
Unsupervised Semantic Discovery Via Global and Local Semantic Alignment in Multimodal Clustering	35248
<i>Zhengzhong Zhu, Pei Zhou, Weihong Du, Shiquan Min, Jiangping Zhu</i>	
ExtendAttack: Attacking Servers of LRMs Via Extending Reasoning.....	35257
<i>Zhenhao Zhu, Yue Liu, Zhiwei Xu, Yingwei Ma, Hongcheng Gao, Nuo Chen, Yanpei Guo, Wenjie Qu, Huiying Xu, Zifeng Kang, Xinzhong Zhu, Jiaheng Zhang</i>	
S ³ -MSD: Large Vision-Language Model for Explainable and Generalizable Multi-Modal Sarcasm Detection	35266
<i>Zhihong Zhu, Fan Zhang, Yunyan Zhang, Jinghan Sun, Guimin Hu, Hao Wu, Yuyan Chen, Bowen Xing, Xian Wu</i>	
CMID: Towards Medical Visual Question Answering Via Contrastive Mutual Information Decoding	35275
<i>Zhihong Zhu, Yunyan Zhang, Fan Zhang, Bowen Xing, Xian Wu</i>	

Response Attack: Exploiting Contextual Priming to Jailbreak Large Language Models 35284
Miao Ziqi, Lijun Li, Yuan Xiong, Zhenhua Liu, Pengyu Zhu, Jing Shao

Ellipsoid-Based Decision Boundaries for Open Intent Classification 35293
Yuetian Zou, Hanlei Zhang, Hua Xu, Songze Li, Long Xiao

Author Index