

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 39: AAAI Technical Tracks

- Natural Language Processing IV

Singapore
20-27 January 2026

ISBN: 979-8-3313-3413-0

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 IV

ProgRAG: Hallucination-Resistant Progressive Retrieval and Reasoning Over Knowledge Graphs	32674
<i>Minbae Park, Hyemin Yang, Jeonghyun Kim, Kunsoo Park, Hyunjoon Kim</i>	
WaterMod: Modular Token-Rank Partitioning for Probability-Balanced LLM Watermarking.....	32683
<i>Shinwoo Park, Hyejin Park, Hyeeseon Ahn, Yo-Sub Han</i>	
FineRef: Fine-Grained Error Reflection and Correction for Long-Form Generation with Citations	32692
<i>Yixing Peng, Licheng Zhang, Shancheng Fang, Yi Liu, Peijian Gu, Quan Wang</i>	
Don't Stop the Multi-Party! on Generating Synthetic Written Multi-Party Conversations with Constraints.....	32701
<i>Nicolò Penzo, Marco Guerini, Bruno Lepri, Goran Glavaš, Sara Tonelli</i>	
GloCTM: Cross-Lingual Topic Modeling Via a Global Context Space	32710
<i>Nguyen Tien Phat, Ngo Vu Minh, Linh Ngo Van, Nguyen Thi Ngoc Diep, Thien Huu Nguyen</i>	
Small Models Exhibit Limited Answer Consistency in Repetition Trials of the Multiple-Choice MMLU-Redux and MedQA Benchmarks	32719
<i>Claudio Pinhanez, Paulo Cavalin, Cassia Sanctos, Marcelo Carpinette Grave</i>	
SOM Directions Are Better than One: Multi-Directional Refusal Suppression in Language Models.....	32728
<i>Giorgio Piras, Raffaele Mura, Fabio Brau, Luca Oneto, Fabio Roli, Battista Biggio</i>	
Facilitating Early Maladaptive Schema-Guided Polite and Empathetic Psychotherapeutic Support: An LLM-Driven MoE-RL-Based Dialogue System.....	32737
<i>Priyanshu Priya, Asif Ekbal</i>	
MirrorShield: Towards Dynamic Adaptive Defense Against Jailbreaks Via Entropy-Guided Mirror Crafting.....	32746
<i>Rui Pu, Chaozhuo Li, Rui Ha, Litian Zhang, Lirong Qiu, Xi Zhang</i>	
MAJIC: Markovian Adaptive Jailbreaking Via Iterative Composition of Diverse Innovative Strategies	32755
<i>Weiwei Qi, Shuo Shao, Wei Gu, Tianhang Zheng, Puning Zhao, Zhan Qin, Kui Ren</i>	
PoeTone: A Framework for Constrained Generation of Structured Chinese Songci with LLMs	32764
<i>Zhan Qu, Shuzhou Yuan, Michael Färber</i>	
RetrySQL: Text-To-SQL Training with Retry Data for Self-Correcting Query Generation.....	32773
<i>Alicja Raczkowska, Riccardo Belluzzo, Piotr Zielinski, Joanna Baran, Pawel Olszewski</i>	
Beyond Fact Retrieval: Episodic Memory for RAG with Generative Semantic Workspaces	32782
<i>Shreyas Rajesh, Pavan Holur, Chenda Duan, David Chong, Vwani Roychowdhury</i>	
Are Language Models Any Good at Density Modeling?.....	32791
<i>Sriram Ranga, Sai Shashank Bedampeta, Rui Mao, Anupam Chattopadhyay</i>	
A Text-Routed Sparse Mixture-Of-Experts Model with Explanation and Temporal Alignment for Multi-Modal Sentiment Analysis.....	32799
<i>Dongning Rao, Yunbiao Zeng, Zhihua Jiang, Jujian Lv</i>	
SPARE: Single-Pass Annotation with Reference-Guided Evaluation for Automatic Process Supervision and Reward Modelling	32808
<i>Md Imbesat Hassan Rizvi, Xiaodan Zhu, Iryna Gurevych</i>	
Anchor Watermark: Robust Attribution for Diffusion-Based Text-To-Audio Model	32817
<i>Xianjin Rong, Donghui Hu</i>	
PASE: Leveraging the Phonological Prior of WavLM for Low-Hallucination Generative Speech Enhancement	32826
<i>Xiaobin Rong, Qinwen Hu, Mansur Yesilbursa, Kamil Wojcicki, Jing Lu</i>	

Speech-Aware Long Context Pruning and Integration for Contextualized Automatic Speech Recognition	32835
<i>Yiming Rong, Yixin Zhang, Ziyi Wang, Deyang Jiang, Yunlong Zhao, Haoran Wu, Shiyu Zhou, Bo Xu</i>	
Interactive Evaluation of Large Language Models for Multi-Requirement Software Engineering Tasks.....	32843
<i>Dimitrios Rontogiannis, Maxime Peyrard, Nicolas Baldwin, Martin Josifoski, Robert West, Dimitrios Gunopulos</i>	
RMO: Towards Better LLM Alignment Via Reshaping Reward Margin Distributions.....	32851
<i>Yanchi Ru, Yue Huang, Xiangliang Zhang</i>	
Incorporating Token Importance in Multi-Vector Retrieval	32860
<i>Archish S, Ankit Garg, Kirankumar Shiragur, Neeraj Kayal</i>	
Assessing the Capabilities of LLMs in Humor: A Multi-Dimensional Analysis of Oogiri Generation and Evaluation.....	32867
<i>Ritsu Sakabe, Hwichan Kim, Toshio Hirasawa, Mamoru Komachi</i>	
Positional Cognitive Specialization: Where Do LLMs Learn to Comprehend and Speak Your Language?	32875
<i>Luis Frentzen Salim, Lun-Wei Ku, Hsing-Kuo Kenneth Pao</i>	
Bonsai: Interpretable Tree-Adaptive Grounded Reasoning.....	32884
<i>Kate Sanders, Benjamin Van Durme</i>	
AntiDote: Bi-Level Adversarial Training for Tamper-Resistant LLMs.....	32893
<i>Debdeep Sanyal, Manodeep Ray, Murari Mandal</i>	
From Chaos to Cure: A Prefix Heuristics Guided Model-Agnostic Adaptive Detoxification Framework.....	32902
<i>Yuhu Shang, Xiang Cheng, Yimeng Ren, Huijia Wu, Xuexiong Luo, Kangkang Lu, Jian Zhao, Zhaofeng He</i>	
Constructing Superior Representations Beyond the Original Documents Via a Contrastive Gaussian Fusion Network for Clustering.....	32911
<i>Ao Shen, Ruizhang Huang, Jingjing Xue, Ruina Bai</i>	
LLMdoctor: Token-Level Flow-Guided Preference Optimization for Efficient Test-Time Alignment of Large Language Models.....	32920
<i>Tiesunlong Shen, Rui Mao, Jin Wang, Heming Sun, Jian Zhang, Xuejie Zhang, Erik Cambria</i>	
Are We on the Right Way to Assess Document Retrieval-Augmented Generation?	32929
<i>Wenxuan Shen, Mingjia Wang, Yaochen Wang, Dongping Chen, Junjie Yang, Yao Wan, Weiwei Lin</i>	
Perturb Your Data: Paraphrase-Guided Training Data Watermarking	32938
<i>Pranav Shetty, Mirazul Haque, Petr Babkin, Zhiqiang Ma, Xiaomo Liu, Manuela Veloso</i>	
Scaling LLM Speculative Decoding: Non-Autoregressive Forecasting in Large-Batch Scenarios.....	32947
<i>Luohe Shi, Zuchao Li, Lefei Zhang, Baoyuan Qi, Guoming Liu, Hai Zhao</i>	
ProFuser: Progressive Fusion of Large Language Models	32956
<i>Tianyuan Shi, Fanqi Wan, Canbin Huang, Xiaojun Quan, Chenliang Li, Ming Yan, Ji Zhang, Minhua Huang, Wu Kai</i>	
From Solver to Tutor: Evaluating the Pedagogical Intelligence of LLMs with KMP-Bench.....	32965
<i>Weikang Shi, Houxing Ren, Junting Pan, Aojun Zhou, Ke Wang, Zimu Lu, Yunqiao Yang, Yuxuan Hu, Linda Wei, Mingjie Zhan, Hongsheng Li</i>	
Enhancing Meme Emotion Understanding with Multi-Level Modality Enhancement and Dual-Stage Modal Fusion.....	32974
<i>Yi Shi, Wenlong Meng, Zhenyuan Guo, Chengkun Wei, Wenzhi Chen</i>	
Fine-Tuned LLMs Know They Don't Know: A Parameter-Efficient Approach to Recovering Honesty.....	32983
<i>Zeyu Shi, Ziming Wang, Tianyu Chen, Shiqi Gao, Haoyi Zhou, Qingyun Sun, Jianxin Li</i>	
Qa-FLoRA: Data-Free Query-Adaptive Fusion of LoRAs for LLMs.....	32992
<i>Shreya Shukla, Aditya Sriram, Milinda Kuppur Narayanaswamy, Hiteshi Jain</i>	

Teaching Large Language Models to Maintain Contextual Faithfulness Via Synthetic Tasks and Reinforcement Learning.....	33001
<i>Shuzheng Si, Haozhe Zhao, Cheng Gao, Yuzhuo Bai, Zhitong Wang, Bofei Gao, Kangyang Luo, Wenhao Li, Yufei Huang, Gang Chen, Fanchao Qi, Minjia Zhang, Baobao Chang, Maosong Sun</i>	
BiCA: Effective Biomedical Dense Retrieval with Citation-Aware Hard Negatives.....	33010
<i>Aarush Sinha, Pavan Kumar S, Roshan Balaji, Nirav Pravinbhai Bhatt</i>	
Causal Reward Adjustment: Mitigating Reward Hacking in External Reasoning Via Backdoor Correction.....	33019
<i>Ruike Song, Zeen Song, Huijie Guo, Wenwen Qiang</i>	
MAVIS: A Benchmark for Multimodal Source Attribution in Long-Form Visual Question Answering.....	33028
<i>Seokwon Song, Minsu Park, Gunhee Kim</i>	
Sparse-DLLM: Accelerating Diffusion LLMs with Dynamic Cache Eviction.....	33038
<i>Yuerong Song, Xiaoran Liu, Ruixiao Li, Zhigeng Liu, Zengfeng Huang, Qipeng Guo, Ziwei He, Xipeng Qiu</i>	
Optimization and Robustness-Informed Membership Inference Attacks for LLMs.....	33047
<i>Zichen Song, Qixin Zhang, Ming Li, Yao Shu</i>	
Chinese Two-Part Allegorical Sayings Reading Comprehension: Exploration from Reasoning to Metaphor.....	33056
<i>Dongyu Su, Yimin Xiao, Tongguan Wang, Feiyue Xue, Junkai Li, Hui Liu, Ying Sha</i>	
CP-Router: An Uncertainty-Aware Router Between LLM and LRM.....	33065
<i>Jiayuan Su, Fulin Lin, Zhaopeng Feng, Han Zheng, Teng Wang, Zhenyu Xiao, Xinlong Zhao, Zuozhu Liu, Lu Cheng, Hongwei Wang</i>	
Beyond Plain Demos: A Demo-Centric Anchoring Paradigm for In-Context Learning in Alzheimer’s Disease Detection.....	33074
<i>Puzhen Su, Haoran Yin, Miao Yongzhu, Jintao Tang, Shasha Li, Ting Wang</i>	
ECD: Evidence-Guided Contrastive Decoding in Retrieval-Augmented Generation with Accurate Knowledge Reference Adjustment.....	33083
<i>Yize Sui, Yan Xu, Kun Hu, Jing Ren, Wenjing Yang</i>	
Bridging the Language Gap: Uncovering and Aligning Shared Circuits for Multi-Hop Reasoning in Multilingual LLMs.....	33091
<i>Chenghao Sun, Zhen Huang, Yonggang Zhang, Xinmei Tian, Xu Shen, Jieping Ye</i>	
Enhancing Pre-Training Data Detection in LLMs Through Discriminative and Symmetric Prefix Selection.....	33100
<i>Kai Sun, Yuxin Lin, Bo Dong, Jingyao Zhang, Bin Shi</i>	
Scope Delineation Before Localization: A Two-Stage Framework for Enhancing Failure Attribution in Multi-Agent Systems.....	33108
<i>Kai Sun, Wenqiang Li, Bo Dong, Yuxin Lin, Jingyao Zhang, Bin Shi</i>	
Efficient Hallucination Detection: Adaptive Bayesian Estimation of Semantic Entropy with Guided Semantic Exploration.....	33117
<i>Qiyao Sun, Xingming Li, Xixiang He, Ao Cheng, Xuanyu Ji, Hailun Lu, Runke Huang, Qingyong Hu</i>	
CO-Bench: Benchmarking Language Model Agents in Algorithm Search for Combinatorial Optimization.....	33126
<i>Weiwei Sun, Shengyu Feng, Shanda Li, Yiming Yang</i>	
Exposing the Cracks: Vulnerabilities of Retrieval-Augmented LLM-Based Machine Translation.....	33135
<i>Yanming Sun, Runzhe Zhan, Chi Seng Cheang, Han Wu, Xuebo Liu, Yuyao Niu, Fengying Ye, Kaixin Lan, Lidia S. Chao, Derek F. Wong</i>	
Well Begun, Half Done: Reinforcement Learning with Prefix Optimization for LLM Reasoning.....	33144
<i>Yiliu Sun, Zicheng Zhao, Yang Wei, Yanfang Zhang, Chen Gong</i>	

CGMIS: Concept-Graph Based Multi-Hop Instructions Synthesis for Enhancing Long-Context Reasoning.....	33153
<i>Zechen Sun, Zecheng Tang, Juntao Li, Wenpeng Hu, Wenliang Chen, Zhunchen Luo, Qiaoming Zhu</i>	
Improving Value-Based Process Verifier Via Low-Cost Variance Reduction.....	33162
<i>Zetian Sun, Dongfang Li, Baotian Hu, Min Zhang</i>	
Improving the Accuracy of Dense Retrieval on the Quantized Indexes Via Gradient Optimization of the Target Embeddings	33171
<i>Cong Tan, Yongqi Shao, Hong Huo, Tao Fang</i>	
PRGB Benchmark: A Robust Placeholder-Assisted Algorithm for Benchmarking Retrieval-Augmented Generation.....	33179
<i>Zhehao Tan, Yihan Jiao, Dan Yang, Junwei Liu, Lei Liu, Jie Feng, Duolin Sun, Yue Shen, Jian Wang, Peng Wei, Jinjie Gu</i>	
RAG-R1: Incentivizing the Search and Reasoning Capabilities of LLMs Through Multi-Query Parallelism	33187
<i>Zhiwen Tan, Jiaming Huang, Qintong Wu, Hongxuan Zhang, Chenyi Zhuang, Jinjie Gu</i>	
Efficient Transcoder Adaptation for Fine-Tuned Models: Revealing Medical Reasoning Mechanisms in Large Language Models	33196
<i>Zhouxing Tan, Hanlin Xue, Yulong Wan, Ruochong Xiong, Xu Chu, Xiang Li, Junfei Liu</i>	
ARQUSUMM: Argument-Aware Quantitative Summarization of Online Conversations	33205
<i>An Quang Tang, Xiuzhen Zhang, Minh Ngoc Dinh, Zhuang Li</i>	
GUI-G?: Gaussian Reward Modeling for GUI Grounding	33214
<i>Fei Tang, Zhangxuan Gu, Zhengxi Lu, Xuyang Liu, Shuheng Shen, Changhua Meng, Wen Wang, Wenqi Zhang, Yongliang Shen, Weiming Lu, Jun Xiao, Yueting Zhuang</i>	
SafetyReminder: Reviving Delayed Safety Awareness of Vision-Language Models to Defend Against Jailbreak Attacks	33223
<i>Peiyuan Tang, Haojie Xin, Xiaodong Zhang, Jun Sun, Qin Xia, Zijiang James Yang</i>	
Put the Space of LoRA Initialization to the Extreme to Preserve Pre-Trained Knowledge.....	33232
<i>Pengwei Tang, Xiaolin Hu, Yong Liu, Lizhong Ding, Dongjie Zhang, Xing Wu, Debing Zhang</i>	
Rectify Evaluation Preference: Improving LLMs' Critique on Math Reasoning Via Perplexity-Aware Reinforcement Learning.....	33241
<i>Changyuan Tian, Zhicong Lu, Shuang Qian, Nayu Liu, Peiguang Li, Li Jin, Leiyi Hu, Zhizhao Zeng, Sirui Wang, Ke Zeng, Guozhi Cas</i>	
PocketLLM: Ultimate Compression of Large Language Models Via Meta Networks.....	33250
<i>Ye Tian, Chengcheng Wang, Jing Han, Yehui Tang, Kai Han</i>	
KeepKV: Achieving Periodic Lossless KV Cache Compression for Efficient LLM Inference.....	33259
<i>Yuxuan Tian, Zihan Wang, Yebo Peng, Aomufei Yuan, Zhiming Wang, Bairen Yi, Xin Liu, Yong Cui, Tong Yang</i>	
PRIME: Planning and Retrieval-Integrated Memory for Enhanced Reasoning	33268
<i>Hieu Tran, Zonghai Yao, Nguyen Luong Tran, Zhichao Yang, Feiyun Ouyang, Shuo Han, Razieh Rahimi, Hong Yu</i>	
Reasoning Transfer for an Extremely Low-Resource and Endangered Language: Bridging Languages Through Sample-Efficient Language Understanding	33277
<i>Khanh-Tung Tran, Barry O'Sullivan, Hoang D. Nguyen</i>	
Listen Like a Teacher: Mitigating Whisper Hallucinations Using Adaptive Layer Attention and Knowledge Distillation.....	33287
<i>Kumud Tripathi, Aditya Srinivas Menon, Aman Gaurav, Raj Prakash Gohil, Pankaj Wasnik</i>	
CharBench: Evaluating the Role of Tokenization in Character-Level Tasks.....	33296
<i>Omri Uzan, Yuval Pinter</i>	
Incoherence as Oracle-Less Measure of Error in LLM-Based Code Generation	33305
<i>Thomas Jean-Michel Valentin, Ardi Madadi, Gaetano Sapia, Marcel Böhme</i>	

Mitigating Content Effects on Reasoning in Language Models Through Fine-Grained Activation Steering.....	33314
<i>Marco Valentino, Geonhee Kim, Dhairya Dalal, Zhixue Zhao, André Freitas</i>	
PRAGWORLD: A Benchmark Evaluating LLMs' Local World Model Under Minimal Linguistic Alterations and Conversational Dynamics	33323
<i>Sachin Vashistha, Aryan Bibhuti, Atharva Naik, Martin Tutek, Somak Aditya</i>	
MCW-KD: Multi-Cost Wasserstein Knowledge Distillation for Large Language Models	33332
<i>Hoang Tran Vuong, Tue Le, Quyen Tran, Linh Ngo Van, Trung Le</i>	
Deep Research Arena: The First Exam of LLMs' Research Abilities Via Seminar-Grounded Tasks.....	33341
<i>Haiyuan Wan, Chen Yang, Junchi Yu, Meiqi Tu, Jiaxuan Lu, Di Yu, Jianbao Cao, Ben Gao, Jiaqing Xie, Aoran Wang, Wenlong Zhang, Philip Torr, Dongzhan Zhou</i>	
Beyond Passive Critical Thinking: Fostering Proactive Questioning to Enhance Human-AI Collaboration	33350
<i>Ante Wang, Yujie Lin, Jingyao Liu, Suhang Wu, Hao Liu, Xinyan Xiao, Jinsong Su</i>	
Deconstructing Pre-Training: Knowledge Attribution Analysis in MoE and Dense Models	33359
<i>Bo Wang, Junzhuo Li, Hong Chen, Yuanlin Chu, Yuxuan Fan, Xuming Hu</i>	
Cross-Granularity Hypergraph Retrieval-Augmented Generation for Multi-Hop Question Answering	33368
<i>Changjian Wang, Weihong Deng, Weili Guan, Quan Lu, Ning Jiang</i>	
Joint Evaluation of Answer and Reasoning Consistency for Hallucination Detection in Large Reasoning Models	33377
<i>Changyue Wang, Weihang Su, Qingyao Ai, Yiqun Liu</i>	
Re-SpS: A Reinforcement Learning Approach to Speculative Sampling	33386
<i>Chenan Wang, Daniel H. Shi, Haipeng Chen</i>	
GRAM-R ² : Self-Training Generative Foundation Reward Models for Reward Reasoning	33395
<i>Chenglong Wang, Yongyu Mu, Hang Zhou, Yifu Huo, Ziming Zhu, Jiali Zeng, Murun Yang, Bei Li, Xiaoyang Hao, Chunliang Zhang, Fandong Meng, Jingbo Zhu, Tong Xiao</i>	
Probing Preference Representations: A Multi-Dimensional Evaluation and Analysis Method for Reward Models.....	33404
<i>Chenglong Wang, Yifu Huo, Yang Gan, Yongyu Mu, Qiaozhi He, Murun Yang, Bei Li, Chunliang Zhang, Tongran Liu, Anxiang Ma, Zhengtao Yu, Jingbo Zhu, Tong Xiao</i>	
ICL-Router: In-Context Learned Model Representations for LLM Routing.....	33413
<i>Chenxu Wang, Hao Li, Yiqun Zhang, Linyao Chen, Jianhao Chen, Ping Jian, Qiaosheng Zhang, Shuyue Hu</i>	
Light-IF: Endowing LLMs with Generalizable Reasoning Via Preview and Self-Checking for Complex Instruction Following.....	33422
<i>Chenyang Wang, Liang Wen, Shousheng Jia, Xiangzheng Zhang, Liang Xu</i>	
Rethinking Flow and Diffusion Bridge Models for Speech Enhancement.....	33431
<i>Dahan Wang, Jun Gao, Tong Lei, Yuxiang Hu, Changbao Zhu, Kai Chen, Jing Lu</i>	
Speech Recognition Model Improves Text-To-Speech Synthesis Using Fine-Grained Reward	33440
<i>Guansu Wang, Peijie Sun</i>	
SDEval: Safety Dynamic Evaluation for Multimodal Large Language Models	33449
<i>Hanqing Wang, Yuan Tian, Mingyu Liu, Zhenhao Zhang, Xiangyang Zhu</i>	
RPTS: Tree-Structured Reasoning Process Scoring for Faithful Multimodal Evaluation	33458
<i>Haofeng Wang, Yu Zhang</i>	
Improving Implicit Discourse Relation Recognition with Natural Language Explanations from LLMs.....	33467
<i>Heng Wang, Changxing Wu</i>	
USE: A Unified Model for Universal Sound Separation and Extraction.....	33476
<i>Hongyu Wang, Chenda Li, Xin Zhou, Shuai Wang, Yanmin Qian</i>	
Enhancing Stability and Fidelity for Zero-Shot TTS with a Multi-Level Evaluator	33485
<i>Hualei Wang, Na Li, Chuke Wang, Shu Wu, Zhifeng Li, Dong Yu</i>	

Bridging the Tokenizer Gap: Semantics and Distribution-Aware Knowledge Transfer for Unbiased Cross-Tokenizer Distillation.....	33494
<i>Huazheng Wang, Yongcheng Jing, Haifeng Sun, Jingyu Wang, Jianxin Liao, Leszek Rutkowski, Dacheng Tao</i>	
A General Highly Accurate Online Planning Method Integrating Large Language Models into Nested Rollout Policy Adaptation for Dialogue Tasks	33503
<i>Hui Wang, Fafa Zhang, Xiaoyu Zhang, Chaoxu Mu</i>	
TTA-Bench: A Comprehensive Benchmark for Evaluating Text-To-Audio Models	33512
<i>Hui Wang, Cheng Liu, Junyang Chen, Haoze Liu, Yuhang Jia, Shiwan Zhao, Jiaming Zhou, Haoqin Sun, Hui Bu, Yong Qin</i>	
ShoppingBench: A Real-World Intent-Grounded Shopping Benchmark for LLM-Based Agents.....	33521
<i>Jiangyuan Wang, Kejun Xiao, Qi Sun, HuaiPeng Zhao, Tao Luo, Jian Dong Zhang, Xiaoyi Zeng</i>	
Accommodate Knowledge Conflicts in Retrieval-Augmented LLMs: Towards Robust Response Generation in the Wild.....	33530
<i>Jiatai Wang, Zhiwei Xu, Di Jin, Xuewen Yang, Tao Li</i>	
Towards Closed-Loop Embodied Empathy Evolution: Probing LLM-Centric Lifelong Empathic Motion Generation in Unseen Scenarios	33539
<i>Jiawen Wang, Jingjing Wang, Tianyang Chen, Min Zhang, Guodong Zhou</i>	
UniCUE: Unified Recognition and Generation Framework for Chinese Cued Speech Video-To-Speech Generation.....	33548
<i>Jinting Wang, Shan Yang, Chenxing Li, Dong Yu, Li Liu</i>	
ComoRAG: A Cognitive-Inspired Memory-Organized RAG for Stateful Long Narrative Reasoning	33557
<i>Juyuan Wang, Rongchen Zhao, Wei Wei, Yufeng Wang, Mo Yu, Jie Zhou, Jin Xu, Liyan Xu</i>	
When Truth is Overridden: Uncovering the Internal Origins of Sycophancy in Large Language Models.....	33566
<i>Keyu Wang, Jin Li, Shu Yang, Zhuoran Zhang, Di Wang</i>	
Decoupling Understanding from Reasoning Via Problem Space Mapping for Small-Scale Model Reasoning	33575
<i>Li Wang, Changhao Zhang, Zengqi Xiu, Kai Lu, Xin Yu, Kui Zhang, Wenjun Wu</i>	
SRACG: A Code Generation Framework with Selective Retrieval Augmentation	33584
<i>Mengzhen Wang, Shukai Ma, Songwen Gong, Jiexin Wang, Ruolin Chen, Liuwen Cao, Yi Cai</i>	
DeepWriter: A Multi-Agent Collaboration Framework for Information-Rich Ultra-Long Book Writing.....	33593
<i>Ming Wang, Minghao Hu, Xiuli Kang, Li He, Yu Tian, Chunming Liu, Han Shi, Zhunchen Luo, Wei Luo, Guotong Geng</i>	
ALEX:A Light Editing-Knowledge Extractor.....	33602
<i>Minghu Wang, Shuliang Zhao, Yuanyuan Zhao, Hongxia Xu</i>	
Beyond Accuracy: A Cognitive Load Framework for Mapping the Capability Boundaries of Tool-Use Agents.....	33611
<i>Qihao Wang, Yue Hu, Mingzhe Lu, Jiayue Wu, Yanbing Liu, Yuanmin Tang</i>	
LoKI: Low-Damage Knowledge Implanting of Large Language Models.....	33620
<i>Runyuan Wang, Peng Ping, Zhengyu Guo, Xiaoye Zhang, Quan Shi, Liting Zhou, Tianbo Ji</i>	

Author Index