

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

Volume 38: AAI Technical Tracks

- Natural Language Processing III

Singapore
20-27 January 2026

ISBN: 979-8-3313-3412-3

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 III

LoopLLM: Transferable Energy-Latency Attacks in LLMs Via Repetitive Generation	31770
<i>Xingyu Li, Xiaolei Liu, Cheng Liu, Yixiao Xu, Kangyi Ding, Bangzhou Xin, Jia-Li Yin</i>	
Do LLMs Feel? Teaching Emotion Recognition with Prompts, Retrieval, and Curriculum Learning	31778
<i>Xinran Li, Yu Liu, Jiaqi Qiao, Xiujuan Xu</i>	
Anti-Adversarial Learning: Desensitizing Prompts for Large Language Model	31787
<i>Xuan Li, Zhe Yin, Xiaodong Gu, Beijun Shen</i>	
VerifyBench: A Systematic Benchmark for Evaluating Reasoning Verifiers Across Domains	31796
<i>Xuzhao Li, Xuchen Li, Shiyu Hu, Yongzhen Guo, Wentao Zhang</i>	
PLaST: Towards Paralinguistic-Aware Speech Translation	31805
<i>Yi Li, Rui Zhao, Ruiquan Zhang, Jinsong Su, Daimeng Wei, Min Zhang, Yidong Chen</i>	
DMOSpeech 2: Reinforcement Learning for Duration Prediction in Metric-Optimized Speech Synthesis	31814
<i>Yinghao Aaron Li, Xilin Jiang, Fei Tao, Cheng Niu, Kaifeng Xu, Juntong Song, Nima Mesgarani</i>	
TransMamba: A Sequence-Level Hybrid Transformer-Mamba Language Model	31823
<i>Yixing Li, Ruobing Xie, Zhen Yang, Xingwu Sun, Shuaipeng Li, Weidong Han, Zhanhui Kang, Di Wang, Yu Cheng</i>	
C ³ TG: Conflict-Aware, Composite, and Collaborative Controlled Text Generation	31834
<i>Yu Li, Zhe Yang, Yi Huang, Xin Liu, Guilin Qi</i>	
AgentSwift: Efficient LLM Agent Design Via Value-Guided Hierarchical Search	31843
<i>Yu Li, Lehui Li, Zhihao Wu, Qingmin Liao, Jianye Hao, Kun Shao, Fengli Xu</i>	
Mitigating Hallucinations in Large Language Models Via Causal Reasoning	31852
<i>Yuangang Li, Yiqing Shen, Yi Nian, Jiechao Gao, Ziyi Wang, Chenxiao Yu, Li Li, Jie Wang, Xiyang Hu, Yue Zhao</i>	
SepPrune: Structured Pruning for Efficient Deep Speech Separation	31861
<i>Yuqi Li, Kai Li, Xin Yin, Zhifei Yang, Zeyu Dong, Zhengtao Yao, Haoyan Xu, Yingli Tian, Yao Lu</i>	
Query-Efficient Domain Knowledge Stealing Against Large Language Models	31870
<i>Zhengao Li, Xiaopeng Yuan, Bolin Shen, Kien Le, Haohan Wang, Xugui Zhou, Shangqian Gao, Yushun Dong</i>	
GraphIF: Enhancing Multi-Turn Instruction Following for Large Language Models with Relation Graph Prompt	31879
<i>Zhenhe Li, Can Lin, Ling Zheng, Wen-Da Wei, Junli Liang, Qi Song</i>	
MacVQA: Adaptive Memory Allocation and Global Noise Filtering for Continual Visual Question Answering	31888
<i>Zhifei Li, Yiran Wang, Chenyi Xiong, Yujing Xia, Xiaoju Hou, Yue Zhao, Miao Zhang, Kui Xiao, Bing Yang</i>	

Efficient Thought Space Exploration Through Strategic Intervention	31897
<i>Ziheng Li, Hengyi Cai, Xiaochi Wei, Yuchen Li, Shuaiqiang Wang, Zhi-Hong Deng, Dawei Yin</i>	
NeSTR: A Neuro-Symbolic Abductive Framework for Temporal Reasoning in Large Language Models.....	31907
<i>Feng Liang, Weixin Zeng, Runhao Zhao, Xiang Zhao</i>	
Bias-Restrained Prefix Representation Finetuning for Mathematical Reasoning.....	31916
<i>Sirui Liang, Pengfei Cao, Jian Zhao, Cong Huang, Jun Zhao, Kang Liu</i>	
Safe RAG by RAG: Untying the Bell that RAG Rang with the RAG Hand.....	31925
<i>Xun Liang, Mengwei Wang, Yuefeng Ma, Simin Niu</i>	
SEAP: Sparse Expert Activation Pruning Unlocks the Brainpower of Large Language Models.....	31934
<i>Xun Liang, Hanyu Wang, Huayi Lai, Simin Niu, Shichao Song, Jiawei Yang, Jihao Zhao, Feiyu Xiong, Bo Tang, Zhiyu Li</i>	
SLD-L2S: Hierarchical Subspace Latent Diffusion for High-Fidelity Lip to Speech Synthesis.....	31943
<i>Yifan Liang, Andong Li, Kang Yang, Guochen Yu, Fangkun Liu, Lingling Dai, Xiaodong Li, Chengshi Zheng</i>	
EquaCode: A Multi-Strategy Jailbreak Approach for Large Language Models Via Equation Solving and Code Completion.....	31952
<i>Zhen Liang, Hai Huang, Zhengkui Chen</i>	
SparK: Query-Aware Unstructured Sparsity with Recoverable KV Cache Channel Pruning	31961
<i>Huanxuan Liao, Yixing Xu, Shizhu He, Guan Chen Li, Xuanwu Yin, Dong Li, Emad Barsoum, Jun Zhao, Kang Liu</i>	
RLMR: Reinforcement Learning with Mixed Rewards for Creative Writing	31970
<i>Jianxing Liao, Tian Zhang, Xiao Feng, Yusong Zhang, Haorui Wang, Bosi Wen, Ziyang Wang, Runzhi Shi</i>	
Query-Routed Activation Editing with Truth-Hierarchical Preference Optimization	31979
<i>Kewei Liao, Tianbo Wang, Yuqing Ma, Zhang Zhang, Zhicheng Geng, Xiaowei Zhao, Jiakai Wang, Xianglong Liu</i>	
DRIFT: Difference-Aware Reinforcement Through Iterative Fine-Tuning for Language Model	31988
<i>Wenjie Liao, Xiaohui Song, Haonan Lu</i>	
GrayKD: Distilling Better Knowledge from Black-Box LLM Via Multi-Rationale Injection.....	31997
<i>Hyeongsoo Lim, Hyung Yong Kim, Jin Young Kim, Min Ho Jang, Eun Seo Seo, Youshin Lim, Shukjae Choi, Jihwan Park, Yunkyu Lim, Hanbin Lee, Byeong-Yeol Kim, Ji Won Yoon</i>	
Towards Better Code Understanding in Decoder-Only Models with Contrastive Learning.....	32006
<i>Jiayi Lin, Yanlin Wang, Yibiao Yang, Lei Zhang, Yutao Xie</i>	
Hidden in the Noise: Unveiling Backdoors in Audio LLMs Alignment Through Latent Acoustic Pattern Triggers	32015
<i>Liang Lin, Miao Yu, Kaiwen Luo, Yibo Zhang, Lilan Peng, Dexian Wang, Xuehai Tang, Yuanhe Zhang, Xikang Yang, Zhenhong Zhou, Kun Wang, Yang Liu</i>	
Diffusion-Assisted Progressive Learning for Weakly Supervised Phrase Localization	32024
<i>Pengyue Lin, Yanyang Hu, Xinjing Liu, Wenqi Jia, Fangxiang Feng, Ruifan Li</i>	

GeWu: A Culturally-Grounded Chinese Benchmark for Multi-Stage Social Bias Evaluation in Large Language Models	32033
<i>Yi Lin, Ziyi Zhou, Jiashi Gao, Xinwei Guo, Jiaxin Zhang, Haiyan Wu, Xin Yao, Xuetao Wei</i>	
RPM-MCTS: Knowledge-Retrieval as Process Reward Model with Monte Carlo Tree Search for Code Generation	32042
<i>Yuanyuan Lin, Xiangyu Ouyang, Teng Zhang, Kaixin Sui</i>	
SparseRM: A Lightweight Preference Modeling with Sparse Autoencoder	32051
<i>Dengcan Liu, Jiahao Li, Zheren Fu, Yi Tu, Jiajun Li, Zhendong Mao, Yongdong Zhang</i>	
Format as a Prior: Quantifying and Analyzing Bias in LLMs for Heterogeneous Data	32060
<i>Jiacheng Liu, Mayi Xu, Qiankun Pi, Wenli Li, Ming Zhong, Yuanyuan Zhu, Mengchi Liu, Tiejun Qian</i>	
Dynamic Cognitive Planning for Cognitive-Functional Dialogue: A Case Study in Emotional Support Conversation	32069
<i>Jiaqi Liu, Yankun Yang, Jiakang Xu, Zhongqiang Du, Wenbin Jiang</i>	
Easy for Children, Hard for AI: The Limits of Multimodal LLMs in Early Childhood Learning	32078
<i>Jingping Liu, Xueyan Wu, Hanxuan Chen, Ziyang Liu, Zhangquan Chen, Ronghao Chen, Huacan Wang</i>	
Is Your (Reasoning) Multimodal Language Model Vulnerable Toward Distractions?	32087
<i>Ming Liu, Hao Chen, Jindong Wang, Liwen Wang, Jingchen Sun, Wensheng Zhang</i>	
Dynamic-Static Synergistic Selection Method for Candidate Code Solutions with Generated Test Cases	32096
<i>Ren-Biao Liu, Jiang-Tian Xue, Chao-Zeng Ma, Hui Sun, Xin-Ye Li, Ming Li</i>	
ARBench: Algorithmic Reasoner Or API Alchemist? Evaluating LLMs Beyond API Calls	32105
<i>Ren-Biao Liu, Chao-Zeng Ma, Anqi Li, Hui Sun, Xin-Ye Li, Ming Li</i>	
Towards Authentic Movie Dubbing with Retrieve-Augmented Director-Actor Interaction Learning	32114
<i>Rui Liu, Yuan Zhao, Zhenqi Jia</i>	
SafeNLIDB: A Privacy-Preserving Safety Alignment Framework for LLM-Based Natural Language Database Interfaces	32123
<i>Ruiheng Liu, Xiaobing Chen, Jinyu Zhang, Qiongwen Zhang, Yu Zhang, Bailong Yang</i>	
MGT-Prism: Enhancing Domain Generalization for Machine-Generated Text Detection Via Spectral Alignment	32132
<i>Shengchao Liu, Xiaoming Liu, Chengzhengxu Li, Zhaohan Zhang, Guoxin Ma, Yu Lan, Shuai Xiao</i>	
BadThink: Triggered Overthinking Attacks on Chain-Of-Thought Reasoning in Large Language Models	32141
<i>Shuaitong Liu, Renjue Li, Lijia Yu, Lijun Zhang, Zhiming Liu, Gaojie Jin</i>	
LLM Collaboration with Multi-Agent Reinforcement Learning	32150
<i>Shuo Liu, Zeyu Liang, Xueguang Lyu, Christopher Amato</i>	
Look as You Think: Unifying Reasoning and Visual Evidence Attribution for Verifiable Document RAG Via Reinforcement Learning	32159
<i>Shuo Chen Liu, Pengfei Luo, Chao Zhang, Yuhao Chen, Haotian Zhang, Qi Liu, Xin Kou, Tong Xu, Enhong Chen</i>	

TruthfulRAG: Resolving Factual-Level Conflicts in Retrieval-Augmented Generation with Knowledge Graphs	32168
<i>Shuyi Liu, Yu-Ming Shang, Xi Zhang</i>	
WaveEx: Accelerating Flow Matching-Based Speech Generation Via Wavelet-Guided Extrapolation	32177
<i>Xiaoqian Liu, Xiyang Gui, Zhengkun Ge, Yuan Ge, Chang Zou, Jiacheng Liu, Zhikang Niu, Qixi Zheng, Chen Xu, Xie Chen, Tong Xiao, Jingbo Zhu, Linfeng Zhang</i>	
LongLLaDA: Unlocking Long Context Capabilities in Diffusion LLMs	32186
<i>Xiaoran Liu, Yuerong Song, Zhigeng Liu, Zengfeng Huang, Qipeng Guo, Ziwei He, Xipeng Qiu</i>	
Focusing on Language: Revealing and Exploiting Language Attention Heads in Multilingual Large Language Models	32195
<i>Xin Liu, Qiyang Song, Qihang Zhou, Haichao Du, Shaowen Xu, Wenbo Jiang, Weijuan Zhang, Xiaoqi Jia</i>	
OX-MABSR: A Benchmark for Open-Domain Explainable Multimodal Aspect-Based Sentiment Reasoning	32204
<i>Xinjing Liu, Zixin Xue, Pengyue Lin, Xinyu Tu, Siwei Xu, Ruifan Li</i>	
Let's Think with Images Efficiently! an Interleaved-Modal Chain-Of-Thought Reasoning Framework with Dynamic and Precise Visual Thoughts	32213
<i>Xu Liu, Yongheng Zhang, Qiguang Chen, Yao Li, Sheng Wang, Libo Qin</i>	
From Detection to Diagnosis: Advancing Hallucination Analysis with Automated Data Synthesis	32222
<i>Yanyi Liu, Qingwen Yang, Tiezheng Guo, Feiyu Qu, Jun Liu, Yingyou Wen</i>	
Answering the Unanswerable is to Err Knowingly: Analyzing and Mitigating Abstention Failures in Large Reasoning Models	32231
<i>Yi Liu, Xiangyu Liu, Zequn Sun, Wei Hu</i>	
Judge Q: Trainable Queries for Optimized Information Retention in KV Cache Eviction	32240
<i>Yijun Liu, Yixuan Wang, Yuzhuang Xu, Shiyu Ji, Yang Xu, Qingfu Zhu, Wanxiang Che</i>	
HEV Generative Sandbox: A Framework for Assessing Domain-Specific Social Risks Through Human-LLM Simulation	32249
<i>Yiran Liu, Zhiyi Hou, Xiaoang Xu, Shuo Wang, Huijia Wu, Kaicheng Yu, Yang Yu, Chengxiang Zhai</i>	
OPERA: A Reinforcement Learning--Enhanced Orchestrated Planner-Executor Architecture for Reasoning-Oriented Multi-Hop Retrieval	32258
<i>Yu Liu, Yanbing Liu, Fangfang Yuan, Cong Cao, Youbang Sun, Kun Peng, Weizhuo Chen, Jianjun Li, Zhiyuan Ma</i>	
InfiGUI-G1: Advancing GUI Grounding with Adaptive Exploration Policy Optimization	32267
<i>Yuhang Liu, Zeyu Liu, Shuanghe Zhu, Pengxiang Li, Congkai Xie, Jiasheng Wang, Xueyu Hu, Xiaotian Han, Jianbo Yuan, Xinyao Wang, Shengyu Zhang, Hongxia Yang, Fei Wu</i>	
Schema-Guided Event Reasoning: A Plug-And-Play Event Reasoning Framework Based on Large Language Models	32276
<i>Yuying Liu, Xuechen Zhao, Yanyi Huang, Ye Wang, Xin Song, Yue Zhang, Haiyan Liu, Bin Zhou</i>	
SlimInfer: Accelerating Long-Context LLM Inference Via Dynamic Token Pruning.....	32284
<i>Lingkun Long, Rubing Yang, Yushi Huang, Desheng Hui, Ao Zhou, Jianlei Yang</i>	

Intuitive Thinking: Expanding Large Language Models' Thinking for Rapid Decision-Making on Candidate Corrections in Chinese Grammar Error Correction.....	32293
<i>Lintao Long, Ruizhang Huang, Ruina Bai, Yongbin Qin, Qihang Fu</i>	
ZipLJP: Zipped Information Processor for Legal Judgment Prediction	32302
<i>Fanghao Lou, Qiqi Wang, Guanyu Chen, Kaiqi Zhao, Huijia Li</i>	
ParaMETA: Towards Learning Disentangled Paralinguistic Speaking Styles Representations from Speech	32311
<i>Haowei Lou, Hye-Young Paik, Wen Hu, Lina Yao</i>	
MCTSr-Zero: Self-Reflective Psychological Counseling Dialogues Generation Via Principles and Adaptive Exploration.....	32320
<i>Hao Lu, Yanchi Gu, Haoyuan Huang, Yulin Zhou, Ningxin Zhu, Chen Li</i>	
URPO: A Unified Reward & Policy Optimization Framework for Large Language Models	32329
<i>Songshuo Lu, Hua Wang, Zhi Chen, Yaohua Tang</i>	
DEALT: LLM-Driven Diversity-Enhanced Data Augmentation for Long-Tail Text Classification.....	32338
<i>Wayne Lu, Xiaoxi Cui</i>	
DomainCQA: Crafting Knowledge-Intensive QA from Domain-Specific Charts.....	32347
<i>Yujing Lu, Ling Zhong, Jing Yang, Weiming Li, Peng Wei, Yongheng Wang, Manni Duan, Qing Zhang</i>	
SpecDetect: Simple, Fast, and Training-Free Detection of LLM-Generated Text Via Spectral Analysis.....	32356
<i>Haitong Luo, Weiyao Zhang, Suhang Wang, Wenji Zou, Chungang Lin, Xuying Meng, Yujun Zhang</i>	
RetroLM: Retrieval-Augmented KVs for Long-Context Processing	32365
<i>Kun Luo, Zheng Liu, Shitao Xiao, Jiabei Chen, Hongjin Qian, Peitian Zhang, Shanshan Jiang, Bin Dong, Jun Zhao, Kang Liu</i>	
Better Datasets Start from RefineLab: Automatic Optimization for High-Quality Dataset Refinement	32374
<i>Xiaonan Luo, Yue Huang, Ping He, Xiangliang Zhang</i>	
Beyond Tokens: Dynamic Latent Reasoning Via Semantic Residual Refinement	32383
<i>Fangrui Lv, Lei Wang, Ruixin Hong, Yong Du, Xiangyu Wu, Tingting Gao, Guorui Zhou, Changshui Zhang</i>	
Trainable EEG Interpolation and Structure-Sharing Dual-Path Encoders for Brain-Assisted Target Speaker Extraction.....	32392
<i>Zhao Lv, Haoran Zhou, Ying Chen, Youdian Gao, Xinhui Li, Ruibo Fu, Cunhang Fan</i>	
A Boundary Token Graph for Zero-Shot Relation Triplet Extraction Involving Discontinuous Entities.....	32401
<i>Kailun Lyu, Zehan Li, Fu Zhang, Jingwei Cheng</i>	
Cueing Without Gapping: Cues-Independent Cued Speech Recognition Powered by Cross-Cues Invariant Modeling	32410
<i>Fengji Ma, Chenxing Li, Li Liu</i>	

ActiShade: Activating Overshadowed Knowledge to Guide Multi-Hop Reasoning in Large Language Models	32419
<i>Huipeng Ma, Luan Zhang, Dandan Song, Linmei Hu, Yuhang Tian, Jun Yang, Changzhi Zhou, Chenhao Li, Yizhou Jin, Xudong Li, Meng Lin, Mingxing Zhang, Shuhao Zhang</i>	
Large Language Models Struggle with Unreasonability in Math Problems	32428
<i>Jingyuan Ma, Damai Dai, Zihang Yuan, Rui Li, Weilin Luo, Bin Wang, Qun Liu, Lei Sha, Zhifang Sui</i>	
Talk2Image: A Multi-Agent System for Multi-Turn Image Generation and Editing.....	32437
<i>Shichao Ma, Yunhe Guo, Jiahao Su, Qihe Huang, Zhengyang Zhou, Yang Wang</i>	
QueryAligner: Customizing User Query to Match LLMs Preferences for Better Intent Recognition	32446
<i>Yunlong Ma, Bo Wang, Yihong Tang, Zifei Yu, Chenyun Xue, Gaoke Zhang, Yuexian Hou</i>	
Inference-Aware Prompt Optimization for Aligning Black-Box Large Language Models	32455
<i>Saaduddin Mahmud, Mason Nakamura, Kyle Hollins Wray, Shlomo Zilberstein</i>	
RefRea: Reference-Guided Reasoning with Meta-Cognition for Accurate Language Model Agents.....	32465
<i>Yuxiang Mai, Qiyue Yin, Wancheng Ni, Jianwei Guo, Xiaogang Ouyang, Pei Xu, Kaiqi Huang</i>	
Confidence Estimation for Text-To-SQL in Large Language Models.....	32474
<i>Sepideh Entezari Maleki, Mohammadreza Pourreza, Davood Rafiei</i>	
BhashaKritika: Building Synthetic Pretraining Data at Scale for Indic Languages	32483
<i>Guduru Manoj, Neel Prabhanjan Rachamalla, Ashish Kulkarni, Gautam Rajeev, Jay Piplodiya, Arul Menezes, Shaharukh Khan, Souvik Rana, Manya Sah, Chandra Khatri, Shubham Agarwal</i>	
MindVote: When AI Meets the Wild West of Social Media Opinion	32492
<i>Xutao Mao, Ezra Xuanru Tao, Leyao Wang</i>	
QuanTaxo: A Quantum Approach to Self-Supervised Taxonomy Expansion.....	32501
<i>Sahil Mishra, Avi Patni, Niladri Chatterjee, Tanmoy Chakraborty</i>	
Textual Self-Attention Network: Test-Time Preference Optimization Through Textual Gradient-Based Attention	32510
<i>Shibing Mo, Haoyang Ruan, Kai Wu, Jing Liu</i>	
Hearing More with Less: Multi-Modal Retrieval-And-Selection Augmented Conversational LLM-Based ASR.....	32519
<i>Bingshen Mu, Hexin Liu, Hongfei Xue, Kun Wei, Lei Xie</i>	
A Multi-Agent LLM Framework for Multi-Domain Low-Resource In-Context NER Via Knowledge Retrieval, Disambiguation and Reflective Analysis	32528
<i>Wenxuan Mu, Jinzhong Ning, Di Zhao, Yijia Zhang</i>	
Learning from the Undesirable: Robust Adaptation of Language Models Without Forgetting	32537
<i>Yunhun Nam, Jaehyung Kim, Jongheon Jeong</i>	
Probabilities Are All You Need: A Probability-Only Approach to Uncertainty Estimation in Large Language Models	32546
<i>Manh Nguyen, Sunil Gupta, Hung Le</i>	
Hallucinate Less by Thinking More: Aspect-Based Causal Abstention for Large Language Models.....	32555
<i>Vy Nguyen, Ziqi Xu, Jeffrey Chan, Estrid He, Feng Xia, Xiuzhen Zhang</i>	

GitTaskBench: A Benchmark for Code Agents Solving Real-World Tasks Through Code Repository Leveraging.....	32564
<i>Ziyi Ni, Huacan Wang, Shuo Zhang, Shuo Lu, Ziyang He, Wangyou , Zhenheng Tang, Sen Hu, Bo Li, Chen Hu, Binxing Jiao, Daxin Jiang, Yuntao Du, Pin Lyu</i>	
Privacy-Protected Retrieval-Augmented Generation for Knowledge Graph Question Answering	32573
<i>Yunfeng Ning, Mayi Xu, Jintao Wen, Qiankun Pi, Yuanyuan Zhu, Ming Zhong, Jiawei Jiang, Tiejun Qian</i>	
TokenPowerBench: Benchmarking the Power Consumption of LLM Inference	32582
<i>Chenxu Niu, Wei Zhang, Jie Li, Yongjian Zhao, Tongyang Wang, Xi Wang, Yong Chen</i>	
Latent Self-Consistency for Reliable Majority-Set Selection in Short- And Long-Answer Reasoning	32591
<i>Jungsuk Oh, Jay-Yoon Lee</i>	
GSAP-ERE: Fine-Grained Scholarly Entity and Relation Extraction Focused on Machine Learning.....	32600
<i>Wolfgang Otto, Lu Gan, Sharmila Upadhyaya, Saurav Karmakar, Stefan Dietze</i>	
GateRA: Token-Aware Modulation for Parameter-Efficient Fine-Tuning	32610
<i>Jie Ou, Shuaihong Jiang, Yingjun Du, Cees G. M. Snoek</i>	
SERL: Self-Examining Reinforcement Learning on Open-Domain	32619
<i>Weixuan Ou, Yanzhao Zheng, Shuoshuo Sun, Wei Zhang, Baohua Dong, Hangcheng Zhu, Ruohui Huang, Gang Yu, Pengwei Yan, Yifan Qiao</i>	
DSCoDeBench: A Realistic Benchmark for Data Science Code Generation	32628
<i>Shuyin Ouyang, Dong Huang, Jingwen Guo, Zeyu Sun, Qihao Zhu, Jie M. Zhang</i>	
Bias Association Discovery Framework for Open-Ended LLM Generations	32637
<i>Jinhao Pan, Chahat Raj, Ziwei Zhu</i>	
Pre-DPO: Improving Data Utilization in Direct Preference Optimization Using a Guiding Reference Model	32646
<i>Junshu Pan, Wei Shen, Shulin Huang, Qiji Zhou, Yue Zhang</i>	
WALKSAFE: Risk-Aware Graph Random Walk with Bi-GRPO for LLM Safety	32655
<i>Shilong Pan, Zhiliang Tian, Wanlong Yu, Zhen Huang, Qingyu Qiu, Zihan Chen, Zhonghao Sun, Minlie Huang, Dongsheng Li</i>	
Explain with Visual Keypoints Like a Real Mentor! a Benchmark for Multimodal Solution Explanation.....	32664
<i>Jaewoo Park, Jungyang Park, Dongju Jang, Jiwan Chung, Byungwoo Yoo, Jaewoo Shin, Seonjoon Park, Taehyeong Kim, Youngjae Yu</i>	

Author Index