

**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 36: AAI Technical Tracks

- Natural Language Processing I

Singapore
20-27 January 2026

ISBN: 979-8-3313-3410-9

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 I

LogicCat: A Chain-Of-Thought Text-To-SQL Benchmark for Complex Reasoning.....	29958
<i>Liutao , Xutao Mao, Dixuan Zhang, Yifan Li, Liuhaixin , Konglulu , Jiaming Hou, Rui Li, Yunlong Li, Aoze Zheng, Zhiqiang Zhang, Luo Zhewei, Hongying Zan, Kunli Zhang, Min Peng</i>	
DialogXpert: Driving Intelligent and Emotion-Aware Conversations Through Online Value-Based Reinforcement Learning with LLM Priors	29967
<i>Tazeek Bin Abdur Rakib, Ambuj Mehrish, Lay-Ki Soon, Wern Han Lim, Soujanya Poria</i>	
Where Norms and References Collide: Evaluating LLMs on Normative Reasoning.....	29976
<i>Mitchell Abrams, Kaveh Eskandari Miandoab, Felix Gervits, Vasanth Sarathy, Matthias Scheutz</i>	
MAMA-Memeia! Multi-Aspect Multi-Agent Collaboration for Depressive Symptoms Identification in Memes.....	29986
<i>Siddhant Agarwal, Adya Dhuler, Polly Ruhnke, Melvin Speisman, Md Shad Akhtar, Shweta Yadav</i>	
ViDia2Std: A Parallel Corpus and Methods for Low-Resource Vietnamese Dialect-To-Standard Translation.....	29995
<i>Khoa Anh Ta, Nguyen Van Dinh, Kiet Van Nguyen</i>	
Rethinking Deep Alignment Through the Lens of Incomplete Safety Learning	30005
<i>Thong Bach, Dung Nguyen, Thao Minh Le, Truyen Tran</i>	
HQ-SVC: Towards High-Quality Zero-Shot Singing Voice Conversion in Low-Resource Scenarios.....	30013
<i>Bingsong Bai, Yizhong Geng, Fengping Wang, Cong Wang, Puyuan Guo, Yingming Gao, Ya Li</i>	
Towards Effective Code-Integrated Reasoning	30022
<i>Fei Bai, Yingqian Min, Beichen Zhang, Zhipeng Chen, Xin Zhao, Lei Fang, Zheng Liu, Zhongyuan Wang, Hongteng Xu</i>	
Identifying and Analyzing Performance-Critical Tokens in Large Language Models.....	30031
<i>Yu Bai, Heyan Huang, Cesare Spinoso-Di Piano, Sanxing Chen, Marc-Antoine Rondeau, Yang Gao, Jackie Chi Kit Cheung</i>	
IROTE: Human-Like Traits Elicitation of Large Language Model Via In-Context Self-Reflective Optimization.....	30040
<i>Yuzhuo Bai, Shitong Duan, Muhua Huang, Jing Yao, Zhenghao Liu, Peng Zhang, Tun Lu, Xiaoyuan Yi, Maosong Sun, Xing Xie</i>	
Time-Frequency Token Advantage Clipping for Training Efficient Large Reasoning Model.....	30049
<i>Rong Bao, Bo Wang, Xiao Wang, Hongyu Li, Rui Zheng, Leszek Rutkowski, Qi Zhang, Liang Ding, Dacheng Tao</i>	
Beyond Next Token Probabilities: Learnable, Fast Detection of Hallucinations and Data Contamination on LLM Output Distributions	30058
<i>Guy Bar-Shalom, Fabrizio Frasca, Derek Lim, Yoav Gelberg, Yftah Ziser, Ran El-Yaniv, Gal Chechik, Haggai Maron</i>	

Steering Pretrained Drafters During Speculative Decoding	30067
<i>Frédéric Berdoz, Peer Rheinboldt, Roger Wattenhofer</i>	
JudgeBoard: Benchmarking and Enhancing Small Language Models for Reasoning Evaluation	30076
<i>Zhenyu Bi, Gaurav Srivastava, Yang Li, Swastik Roy, Meng Lu, Morteza Ziyadi, Xuan Wang</i>	
OptiHive: Ensemble Selection for LLM-Based Optimization Via Statistical Modeling	30085
<i>Maxime Bouscary, Saurabh Amin</i>	
Do LLMs Really Struggle at NL-FOL Translation? Revealing Their Strengths Via a Novel Benchmarking Strategy	30094
<i>Andrea Brunello, Luca Geatti, Michele Mignani, Angelo Montanari, Nicola Saccomanno</i>	
Sampling-Free Uncertainty Quantification Via Hidden State Dynamics in Language Models	30104
<i>Yixin Bu, Guanyun Zou, Renzhi Wang, Runze Xia, Cunjun Wang, Hongliang Dai, Xiaoqing Ma, Piji Li</i>	
RaCoT: Plug-And-Play Contrastive Example Generation Mechanism for Enhanced LLM Reasoning Reliability	30112
<i>Kaitong Cai, Jusheng Zhang, Yijia Fan, Jing Yang, Keze Wang</i>	
ERank: Fusing Supervised Fine-Tuning and Reinforcement Learning for Effective and Efficient Text Reranking	30121
<i>Yuzheng Cai, Yanzhao Zhang, Dingkun Long, Mingxin Li, Pengjun Xie, Weiguo Zheng</i>	
Does Question Really Matter? the Attribution of Answer Bias in LLM Evaluation	30130
<i>Boxi Cao, Ruotong Pan, Hongyu Lin, Xianpei Han, Le Sun</i>	
On the Superimposed Noise Accumulation Problem in Sequential Knowledge Editing of Large Language Models	30139
<i>Ding Cao, Yuchen Cai, Yuqing Huang, Xuesong He, Rongxi Guo, Guiquan Liu, Guangzhong Sun</i>	
TIV: Thought Injection Via Vectors for Efficient Reasoning in Large Reasoning Models.....	30148
<i>Yi Cao, Weijie Shi, Wei-Jie Xu, Yucheng Shen, Yue Cui, Hanghui Guo, Shimin Di, Ziyi Liu, Jiaming Li, Alexander Zhou, Jia Zhu, Jiajie Xu</i>	
DCTR: Dual-Constraint Subgraph Optimization for Knowledge Graph-Based Retrieval- Augmented Generation.....	30156
<i>Yukun Cao, Zirui Xu, Dongyang Li, Zhihao Guo, Luobin Huang, Lisheng Wang</i>	
Sparse Attention Across Multiple-Context KV Cache.....	30165
<i>Ziyi Cao, Qingyi Si, Jingbin Zhang, Bingquan Liu</i>	
Calibrating and Rotating: A Unified Framework for Weight Conditioning in PEFT	30174
<i>Da Chang, Peng Xue, Yu Li, Yongxiang Liu, Pengxiang Xu, Shixun Zhang</i>	
PathRAG: Pruning Graph-Based Retrieval Augmented Generation with Relational Paths	30183
<i>Boyu Chen, Zirui Guo, Zidan Yang, Yuluo Chen, Junze Chen, Zhenghao Liu, Chuan Shi, Cheng Yang</i>	
Can Editing LLMs Inject Harm?	30192
<i>Canyu Chen, Baixiang Huang, Zekun Li, Zhaorun Chen, Shiyang Lai, Xiong Xiao Xu, Jia-Chen Gu, Jindong Gu, Huaxiu Yao, Chaowei Xiao, Xifeng Yan, William Yang Wang, Philip Torr, Dawn Song, Kai Shu</i>	

LiR3AG: A Lightweight Rerank Reasoning Strategy Framework for Retrieval-Augmented Generation	30201
<i>Guo Chen, Junjie Huang, Huaijin Xie, Fei Sun, Tao Jia</i>	
Skill Path: Unveiling Language Skills from Circuit Graphs	30210
<i>Hang Chen, Xinyu Yang, Jiaying Zhu, Wenya Wang</i>	
AUVIC: Adversarial Unlearning of Visual Concepts for Multi-Modal Large Language Models	30218
<i>Haokun Chen, Jianing Li, Yao Zhang, Jinhe Bi, Yan Xia, Jindong Gu, Volker Tresp</i>	
Rethinking Label Consistency of In-Context Learning: An Implicit Transductive Label Propagation Perspective.....	30226
<i>Haoyang Chen, Richong Zhang, Junfan Chen</i>	
Auto-PRE: An Automatic and Cost-Efficient Peer-Review Framework for Language Generation Evaluation.....	30235
<i>Junjie Chen, Weihang Su, Zhumin Chu, Haitao Li, Yujia Zhou, Dingbo Yuan, Xudong Wang, Jun Zhou, Yiqun Liu, Min Zhang, Shaoping Ma, Qingyao Ai</i>	
Learning from Scoring Disagreements: Contrastive Error Mining for Efficient and Robust LLM-Based Assessment.....	30243
<i>Lei Chen, Tengting Cheng, Boyu Gao, Zitao Liu, Weiqi Luo</i>	
DeFuzzRAG: Handling Fuzzy Time Expressions for Temporal Robustness in Retrieval-Augmented Generation	30252
<i>Ling-Chun Chen, Hsi-Wen Chen, Ming-Syan Chen</i>	
Aware First, Think Less: Dynamic Boundary Self-Awareness Drives Significant Gains in Reasoning Efficiency in Large Language Models.....	30261
<i>Qiguang Chen, Dengyun Peng, Jinhao Liu, Huikang Su, Jiannan Guan, Libo Qin, Wanxiang Che</i>	
You Don't Need Pre-Built Graphs for RAG: Retrieval Augmented Generation with Adaptive Reasoning Structures	30270
<i>Shengyuan Chen, Chuang Zhou, Zheng Yuan, Qinggang Zhang, Zeyang Cui, Hao Chen, Yilin Xiao, Jiannong Cao, Xiao Huang</i>	
DEPO: Dual-Efficiency Preference Optimization for LLM Agents	30279
<i>Sirui Chen, Mengshi Zhao, Lei Xu, Yuying Zhao, Beier Zhu, Hanwang Zhang, Shengjie Zhao, Chaochao Lu</i>	
Flora: Effortless Context Construction to Arbitrary Length and Scale.....	30288
<i>Tianxiang Chen, Zhentao Tan, Xiaofan Bo, Yue Wu, Tao Gong, Qi Chu, Jieping Ye</i>	
How Does Chain of Thought Think? Mechanistic Interpretability of Chain-Of-Thought Reasoning with Sparse Autoencoding.....	30297
<i>Xi Chen, Aske Plaat, Niki Van Stein</i>	
ModalSyncSum: Synchronizing Image and Text for Reliable Summary Generation.....	30306
<i>Xuanqi Chen, Ziyang Rong, Xinfeng Liao, Yiqian Wu, Bowei Zhang, Pengfei Fu, Shengyi Jiang</i>	
Improving Long-Context Summarization with Multi-Granularity Retrieval Optimization.....	30315
<i>Xueyu Chen, Kaitao Song, Zifan Song, Dongsheng Li, Cairong Zhao</i>	
A Better Start: Sensitivity-Aware Warm-Up for Robust and Efficient Fine-Tuning.....	30324
<i>Yile Chen, Zeyi Wen, Jian Chen, Jin Huang</i>	

Schema-Guided Scene-Graph Reasoning Based on Multi-Agent Large Language Model System	30332
<i>Yiye Chen, Harpreet S. Sawhney, Nicholas Gydé, Yanan Jian, Jack Saunders, Patricio Vela, Benjamin E Lundell</i>	
DiagramGPT-Llama3:Enabling Editable, High-Fidelity Diagram Generation with Vision Large Language Models	30341
<i>Yongyuan Chen, Minjie Hong, Boxi Wu, Xicheng Han</i>	
CounterBench: Evaluating and Improving Counterfactual Reasoning in Large Language Models	30350
<i>Yuefei Chen, Vivek K. Singh, Jing Ma, Ruixiang Tang</i>	
StoryBox: Collaborative Multi-Agent Simulation for Hybrid Bottom-Up Long-Form Story Generation Using Large Language Models	30359
<i>Zehao Chen, Rong Pan, Haoran Li</i>	
From Mathematical Reasoning to Code: Generalization of Process Reward Models in Test-Time Scaling	30368
<i>Zhengyu Chen, Yudong Wang, Teng Xiao, Ruochen Zhou, Xusheng Yang, Wei Wang, Zhifang Sui, Jingang Wang</i>	
Reasoning with Exploration: An Entropy Perspective	30377
<i>Daixuan Cheng, Shaohan Huang, Xuekai Zhu, Bo Dai, Xin Zhao, Zhenliang Zhang, Furu Wei</i>	
Human Cognition Inspired RAG with Knowledge Graph for Complex Problem Solving.....	30386
<i>Yao Cheng, Yibo Zhao, Jiapeng Zhu, Yao Liu, Xing Sun, Xiang Li</i>	
Activations as Features: Probing LLMs for Generalizable Essay Scoring Representations.....	30395
<i>Jinwei Chi, Ke Wang, Yu Chen, Xuanye Lin, Qiang Xu</i>	
“As Eastern Powers, I Will Veto.”: An Investigation of Nation-Level Bias of Large Language Models in International Relations.....	30404
<i>Jonghyeon Choi, Yeonjun Choi, Hyun-Chul Kim, Beakcheol Jang</i>	
HanjaBridge: Resolving Semantic Ambiguity in Korean LLMs Via Hanja-Augmented Pre-Training	30413
<i>Seungho Choi, Sihyun Park, Minsang Kim, Chansol Park, Bongsu Kim</i>	
Persistent Backdoor Attacks Under Continual Fine-Tuning of LLMs.....	30422
<i>Jing Cui, Yufei Han, Jianbin Jiao, Junge Zhang</i>	
When Smiley Turns Hostile: Interpreting How Emojis Trigger LLMs’ Toxicity	30431
<i>Shiyao Cui, Xijia Feng, Yingkang Wang, Junxiao Yang, Zhixin Zhang, Biplab Sikdar, Hongning Wang, Han Qiu, Minlie Huang</i>	
HLPD: Aligning LLMs to Human Language Preference for Machine-Revised Text Detection	30440
<i>Fangqi Dai, Xingjian Jiang, Zizhuang Deng</i>	
GOMPSNR: Reflourish the Signal-To-Noise Ratio Metric for Audio Generation Tasks.....	30449
<i>Lingling Dai, Andong Li, Cheng Chi, Yifan Liang, Xiaodong Li, Chengshi Zheng</i>	
Careful Queries, Credible Results: Teaching RAG Models Advanced Web Search Tools with Reinforcement Learning	30458
<i>Yuqin Dai, Shuo Yang, Guoqing Wang, Yong Deng, Zhanwei Zhang, Jun Yin, Pengyu Zeng, Zhenzhe Ying, Changhua Meng, Can Yi, Yuchen Zhou, Weiqiang Wang, Shuai Lu</i>	

Learner-Tailored Program Repair: A Solution Generator with Iterative Edit-Driven Retrieval Enhancement	30467
<i>Zhenlong Dai, Zhuoluo Zhao, Hengning Wang, Xiu Tang, Sai Wu, Chang Yao, Zhipeng Gao, Jingyuan Chen</i>	
Breaking the Adversarial Robustness-Performance Trade-Off in Text Classification Via Manifold Purification	30476
<i>Chenhao Dang, Jing Ma</i>	
Measuring the Unmeasurable: Unveiling Latent Cognitive Capabilities of LLM.....	30485
<i>Cui Danxin, Sihang Jiang, Keyi Wang, Zhiyi Duan, Yanghua Xiao, Bi Yude, Jiaqing Liang, Mingguo He, Shimin Tao, Yilun Liu</i>	
CLER: Improving Multimodal Financial Reasoning by Cross-MLLM Error Reflection	30494
<i>Shuangyan Deng, Zhongsheng Wang, Rui Mao, Ciprian Doru Giurcaneanu, Jiamou Liu</i>	
Guess Or Recall? Training CNNs to Classify and Localize Memorization in LLMs.....	30503
<i>J�r�mie Dentan, Davide Buscaldi, Sonia Vanier</i>	
SkillGen: Learning Domain Skills for In-Context Sequential Decision Making	30512
<i>Ruomeng Ding, Wei Cheng, Minglai Shao, Chen Zhao</i>	
W2S-AlignTree: Weak-To-Strong Inference-Time Alignment for Large Language Models Via Monte Carlo Tree Search.....	30521
<i>Zhenyu Ding, Yuhao Wang, Tengyue Xiao, Haoying Wang, Guojun Ma, Mingyang Wan, Caigui Jiang, Ning Ding</i>	
AdaSpec: Adaptive Multilingual Speculative Decoding with Self-Synthesized Language-Aware Training and Vocabulary Simplification	30530
<i>Dinh-Truong Do, Nguyen-Khang Le, Le-Minh Nguyen</i>	
SACodec: Asymmetric Quantization with Semantic Anchoring for Low-Bitrate High-Fidelity Neural Speech Codecs.....	30539
<i>Zhongren Dong, Bin Wang, Jing Han, Haotian Guo, Xiaojun Mo, Yimin Cao, Zixing Zhang</i>	
TO-GATE: Clarifying Questions and Summarizing Responses with Trajectory Optimization for Eliciting Human Preference.....	30548
<i>Yulin Dou, Jiangming Liu</i>	
Mnemosyne: Accelerating Multi-Hop Question Answering Via Cache Hit Order Fitting	30557
<i>Haizhou Du, Jiujiu Li, Dongyang Li, Luobin Huang, Lisheng Wang</i>	
Scaling Towards the Information Boundary of Instructions Through Data Synthesizing	30566
<i>Li Du, Hanyu Zhao, Yiming Ju, Tengfei Pan</i>	
CCFQA: A Benchmark for Cross-Lingual and Cross-Modal Speech and Text Factuality Evaluation.....	30575
<i>Yexing Du, Kaiyuan Liu, Youcheng Pan, Zheng Chu, Bo Yang, Xiaocheng Feng, Ming Liu, Yang Xiang</i>	
MemGuide: Intent-Driven Memory Selection for Goal-Oriented Multi-Session LLM Agents	30584
<i>Yiming Du, Bingbing Wang, Yang He, Bin Liang, Baojun Wang, Zhongyang Li, Lin Gui, Jeff Z. Pan, Ruifeng Xu, Kam-Fai Wong</i>	
Test-Time Reinforcement Learning for GUI Grounding Via Region Consistency.....	30593
<i>Yong Du, Yuchen Yan, Fei Tang, Zhengxi Lu, Chang Zong, Weiming Lu, Shengpei Jiang, Yongliang Shen</i>	

Learning to Parse and Reconstruct: Bidirectional Modeling of Question-To-Program Mapping.....	30602
<i>Zeying Duan, Youtian Du, Yuanlin Chang, Bowen Lin, Weijia Wu</i>	
HC2-GNN: Hierarchical Graph Representation Learning for Efficient Text Classification.....	30611
<i>Jiejie Fan, Xiaojuan Ban, Zhiyan Zhang, Xi Sun</i>	
TimeBill: Time-Budgeted Inference for Large Language Models.....	30620
<i>Qi Fan, An Zou, Yehan Ma</i>	
What Makes a Good Speech Tokenizer for LLM-Centric Speech Generation? a Systematic Study.....	30629
<i>Xiaoran Fan, Zhichao Sun, Yangfan Gao, Jingfei Xiong, Hang Yan, Yifei Cao, Jiajun Sun, Shuo Li, Zhihao Zhang, Zhiheng Xi, Yuhao Zhou, Senjie Jin, Changhao Jiang, Junjie Ye, Ming Zhang, Rui Zheng, Zhenhua Han, Yunke Zhang, Demei Yan, Shaokang Dong, Tao Ji, Tao Gui</i>	
MHA2MLA-VLM: Enabling DeepSeek’s Economical Multi-Head Latent Attention Across Vision-Language Models	30638
<i>Xiaoran Fan, Zhichao Sun, Tao Ji, Lixing Shen, Tao Gui</i>	
SegMem-RAG: Adaptive Memory for Retrieval-Augmented Generation in Open-Ended Knowledge Environments	30647
<i>Xuanbo Fan, Tianqi Zhao, Yi Cheng, Chi Xiu, Jiaxin Guo, Boci Peng, Bingjing Xu, Jessica Zhang, Feng Sun, Yan Zhang</i>	
GlitchCleaner: Lightweight Glitch Tokens Repairing by Lossless Gated LoRA in Large Language Models.....	30656
<i>Yibo Fan, Jingru Li, Huan Li</i>	
Graph of Verification: Structured Verification of LLM Reasoning with Directed Acyclic Graphs	30665
<i>Jiwei Fang, Bin Zhang, Changwei Wang, Jin Wan, Zhiwei Xu</i>	
Boosting ASR Robustness Via Test-Time Reinforcement Learning with Audio-Text Semantic Rewards.....	30673
<i>Linghan Fang, Tianxin Xie, Li Liu</i>	
RESTL: Reinforcement Learning Guided by Multi-Aspect Rewards for Signal Temporal Logic Transformation	30682
<i>Yue Fang, Zhi Jin, Jie An, Hongshen Chen, Xiaohong Chen, Naijun Zhan</i>	
Toward Better EHR Reasoning in LLMs: Reinforcement Learning with Expert Attention Guidance.....	30690
<i>Yue Fang, Yuxin Guo, Jiaran Gao, Hongxin Ding, Xinke Jiang, Weibin Liao, Yongxin Xu, Yinghao Zhu, Zhibang Yang, Liantao Ma, Junfeng Zhao, Yasha Wang</i>	
SCIR: A Self-Correcting Iterative Refinement Framework for Enhanced Information Extraction Based on Schema.....	30699
<i>Yushen Fang, Jianjun Li, Mingqian Ding, Chang Liu, Xinchu Zou, Wenqi Yang</i>	
Towards Better Correctness and Efficiency in Code Generation.....	30708
<i>Yunlong Feng, Yang Xu, Xiao Xu, Binyuan Hui, Junyang Lin</i>	
PASS: Probabilistic Agentic Supernet Sampling for Interpretable and Adaptive Chest X-Ray Reasoning	30717
<i>Yushi Feng, Junye Du, Yingying Hong, Qifan Wang, Lequan Yu</i>	
Graph-Of-Mark: Promote Spatial Reasoning in Multimodal Language Models with Graph-Based Visual Prompting	30726
<i>Giacomo Frisoni, Lorenzo Molletta, Mattia Buzzoni, Gianluca Moro</i>	

Efficient and Adaptive Simultaneous Speech Translation with Fully Unidirectional Architecture	30735
<i>Biao Fu, Donglei Yu, Minpeng Liao, Chengxi Li, Xinjie Chen, Yidong Chen, Kai Fan, Xiaodong Shi</i>	
The Visual Prism: Refracting Images into Parallel Multilingual Descriptions with Structured Visual Guidance.....	30744
<i>Chengpeng Fu, Xiaocheng Feng, Yichong Huang, Wenshuai Huo, Baohang Li, Yang Xiang, Ting Liu</i>	
GraphIC: A Graph-Based In-Context Example Retrieval Model for Multi-Step Reasoning.....	30753
<i>Jiale Fu, Yaqing Wang, Simeng Han, Jiaming Fan, Xu Yang</i>	
Harnessing the Unseen: The Hidden Influence of Intrinsic Knowledge in Long-Context Language Models.....	30762
<i>Yu Fu, Haz Sameen Shahgir, Hui Liu, Xianfeng Tang, Qi He, Yue Dong</i>	
I Have Covered All the Bases Here: Interpreting Reasoning Features in Large Language Models Via Sparse Autoencoders.....	30771
<i>Andrey V. Galichin, Alexey Dontsov, Polina Druzhinina, Anton Razzhigaev, Oleg Rogov, Elena Tutubalina, Ivan Oseledets</i>	
S2D-Align: Shallow-To-Deep Auxiliary Learning for Anatomically-Grounded Radiology Report Generation	30780
<i>Jiechao Gao, Chang Liu, Yuangang Li</i>	
DeCoRL: Decoupling Reasoning Chains Via Parallel Sub-Step Generation and Cascaded Reinforcement for Interpretable and Scalable RLHF	30789
<i>Ziyuan Gao, Di Liang, Xianjie Wu, Philippe Morel, Minlong Peng</i>	
Expert-Guided Prompting and Retrieval-Augmented Generation for Emergency Medical Service Question Answering	30798
<i>Xueren Ge, Sahil Murtaza, Anthony Cortez, Homa Alemzadeh</i>	
SageLM: A Multi-Aspect and Explainable Large Language Model for Speech Judgement	30807
<i>Yuan Ge, Junxiang Zhang, Xiaoqian Liu, Bei Li, Xiangnan Ma, Chenglong Wang, Kaiyang Ye, Yangfan Du, Linfeng Zhang, Yuxin Huang, Tong Xiao, Zhengtao Yu, Jingbo Zhu</i>	
Control Illusion: The Failure of Instruction Hierarchies in Large Language Models.....	30816
<i>Yilin Geng, Haonan Li, Honglin Mu, Xudong Han, Timothy Baldwin, Omri Abend, Eduard Hovy, Lea Frermann</i>	
Analyze–Compose–Execute: A Dynamic Dialogue Framework for Multi-Agent Debate	30825
<i>Wenyuan Gu, Haowen Wang, Jiale Han, Xiang Li, Zhixuan Wu, Hongru Xiao, Bo Cheng</i>	
Group Causal Policy Optimization for Post-Training Large Language Models.....	30834
<i>Ziyin Gu, Jingyao Wang, Ran Zuo, Chuxiong Sun, Zeen Song, Changwen Zheng, Wenwen Qiang</i>	
Beware of Reasoning Overconfidence: Pitfalls in the Reasoning Process for Multi-Solution Tasks	30843
<i>Jiannan Guan, Qiguang Chen, Libo Qin, Dengyun Peng, Jinhao Liu, Liangyu Huo, Jian Xie, Wanxiang Che</i>	
Beyond Superficial Forgetting: Thorough Unlearning Through Knowledge Density Estimation and Block Re-Insertion	30852
<i>Feng Guo, Yuntao Wen, Shen Gao, Junshuo Zhang, Shuo Shang</i>	

Author Index