

39th AAAI Conference on Artificial Intelligence (AAAI-25), 37th Conference on Innovative Applications of Artificial Intelligence (IAAI-25), and 15th Symposium on Educational Advances in Artificial Intelligence (EAAI-25)

Volume 24: AAAI Technical Tracks

- Natural Language Processing III

Philadelphia, Pennsylvania, USA
25 February - 4 March 2025

Part 1 of 2

ISBN: 979-8-3313-2534-3

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© (2025) by Association for the Advancement of Artificial Intelligence
All rights reserved.

Printed with permission by Curran Associates, Inc. (2025)

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

PART 1

AAAI TECHNICAL TRACK ON NATURAL LANGUAGE PROCESSING III

Is Your Image a Good Storyteller?	25165
<i>Xiujie Song, Xiaoyi Pang, Haifeng Tang, Mengyue Wu, Kenny Q. Zhu</i>	
ELLA-V: Stable Neural Codec Language Modeling with Alignment-Guided Sequence Reordering.....	25174
<i>Yakun Song, Zhuo Chen, Xiaofei Wang, Ziyang Ma, Xie Chen</i>	
Tradutor: Building a Variety Specific Translation Model.....	25183
<i>Hugo Sousa, Satya Almasian, Ricardo Campos, Alipio Jorge</i>	
Enhancing Portuguese Variety Identification with Cross-Domain Approaches.....	25192
<i>Hugo Sousa, Rúben Almeida, Purificação Silvano, Inês Cantante, Ricardo Campos, Alipio Jorge</i>	
Structured Packing in LLM Training Improves Long Context Utilization.....	25201
<i>Konrad Staniszewski, Szymon Tworowski, Sebastian Jaszczur, Yu Zhao, Henryk Michalewski, Łukasz Kuciński, Piotr Miłoś</i>	
VERO: Verification and Zero-Shot Feedback Acquisition for Few-Shot Multimodal Aspect-Level Sentiment Classification.....	25210
<i>Kai Sun, Hao Wu, Bin Shi, Samuel Mensah, Peng Liu, Bo Dong</i>	
Transformer Layers as Painters	25219
<i>Qi Sun, Marc Pickett, Aakash Kumar Nain, Llion Jones</i>	
ExDDI: Explaining Drug-Drug Interaction Predictions with Natural Language.....	25228
<i>Zhaoyue Sun, Jiazheng Li, Gabriele Pergola, Yulan He</i>	
Tuning-Free Accountable Intervention for LLM Deployment – a Metacognitive Approach	25237
<i>Zhen Tan, Jie Peng, Song Wang, Lijie Hu, Tianlong Chen, Huan Liu</i>	
Multi-Grained Query-Guided Set Prediction Network for Grounded Multimodal Named Entity Recognition	25246
<i>Jielong Tang, Zhenxing Wang, ZiYang Gong, Jianxing Yu, Xiangwei Zhu, Jian Yin</i>	
Adaptive Few-shot Prompting for Machine Translation with Pre-trained Language Models	25255
<i>Lei Tang, Jinghui Qin, Wenxuan Ye, Hao Tan, Zhijing Yang</i>	
Unleashing the Potential of Large Language Models as Prompt Optimizers: Analogical Analysis with Gradient-based Model Optimizers.....	25264
<i>Xinyu Tang, Xiaolei Wang, Wayne Xin Zhao, Siyuan Lu, Yaliang Li, Ji-Rong Wen</i>	
A Comprehensive Evaluation on Event Reasoning of Large Language Models	25273
<i>Zhengwei Tao, Zhi Jin, Yifan Zhang, Xiancai Chen, Haiyan Zhao, Jia Li, Bin Liang, Chongyang Tao, Qun Liu, Kam-Fai Wong</i>	
Few-Shot, No Problem: Descriptive Continual Relation Extraction	25282
<i>Nguyen Xuan Thanh, Anh Duc Le, Quyen Tran, Thanh-Thien Le, Linh Ngo Van, Thien Huu Nguyen</i>	

A Systematic Exploration of Knowledge Graph Alignment with Large Language Models in Retrieval Augmented Generation	25291
<i>Shiyu Tian, Shuyue Xing, Xingrui Li, Yangyang Luo, Caixia Yuan, Wei Chen, Huixing Jiang, Xiaojie Wang</i>	
CodeHalu: Investigating Code Hallucinations in LLMs via Execution-based Verification	25300
<i>Yuchen Tian, Weixiang Yan, Qian Yang, Xuandong Zhao, Qian Chen, Wen Wang, Ziyang Luo, Lei Ma, Dawn Song</i>	
Infer Human's Intentions Before Following Natural Language Instructions.....	25309
<i>Yanming Wan, Yue Wu, Yiping Wang, Jiayuan Mao, Natasha Jaques</i>	
LiteSearch: Efficient Tree Search with Dynamic Exploration Budget for Math Reasoning	25318
<i>Ante Wang, Linfeng Song, Ye Tian, Baolin Peng, Dian Yu, Haitao Mi, Jinsong Su, Dong Yu</i>	
A New Formula for Sticker Retrieval: Reply with Stickers in Multi-Modal and Multi-Session Conversation.....	25327
<i>Bingbing Wang, Yiming Du, Bin Liang, Zhixin Bai, Min Yang, Baojun Wang, Kam-Fai Wong, Rui Feng Xu</i>	
RoVRM: A Robust Visual Reward Model Optimized via Auxiliary Textual Preference Data	25336
<i>Chenglong Wang, Yang Gan, Yifu Huo, Yongyu Mu, Murun Yang, Qiaozhi He, Tong Xiao, Chunliang Zhang, Tongran Liu, Jingbo Zhu</i>	
Thought-Path Contrastive Learning via Premise-Oriented Data Augmentation for Logical Reading Comprehension.....	25345
<i>Chenxu Wang, Ping Jian, Zhen Yang</i>	
Decoupling Metacognition from Cognition: A Framework for Quantifying Metacognitive Ability in LLMs.....	25353
<i>Guoqing Wang, Wen Wu, Guangze Ye, Zhenxiao Cheng, Xi Chen, Hong Zheng</i>	
Self-Evolutionary Large Language Models Through Uncertainty-Enhanced Preference Optimization.....	25362
<i>Jianing Wang, Yang Zhou, Xiaocheng Zhang, Mengjiao Bao, Peng Yan</i>	
STAMPsy: Towards SpatioTemporal-Aware Mixed-Type Dialogues for Psychological Counseling	25371
<i>Jieyi Wang, Yue Huang, Zeming Liu, Dexuan Xu, Chuan Wang, Xiaoming Shi, Ruiyuan Guan, Hongxing Wang, Weihua Yue, Yu Huang</i>	
SECodec: Structural Entropy-based Compressive Speech Representation Codec for Speech Language Models	25380
<i>Linqin Wang, Yaping Liu, Zhengtao Yu, Shengxiang Gao, Cunli Mao, Yuxin Huang, Wenjun Wang, Ling Dong</i>	
Distilling Structured Rationale from Large Language Models to Small Language Models for Abstractive Summarization	25389
<i>Linyong Wang, Lianwei Wu, Shaoqi Song, Yaxiong Wang, Cuiyun Gao, Kang Wang</i>	
McHirc: A Multimodal Benchmark for Chinese Idiom Reading Comprehension.....	25398
<i>Tongguan Wang, Mingmin Wu, Guixin Su, Dongyu Su, Yuxue Hu, Zhongqiang Huang, Ying Sha</i>	
Int*-Match: Balancing Intra-Class Compactness and Inter-Class Discrepancy for Semi-Supervised Speaker Recognition.....	25407
<i>Xingmei Wang, Jinghan Liu, Jiayang Meng, Boquan Li, Zijian Liu</i>	

Exploring Activation Patterns of Parameters in Language Models	25416
<i>Yudong Wang, Damai Dai, Zhe Yang, Jingyuan Ma, Zhifang Sui</i>	
Friends-MMC: A Dataset for Multi-modal Multi-party Conversation Understanding	25425
<i>Yueqian Wang, Xiaojun Meng, Yuxuan Wang, Jianxin Liang, Qun Liu, Dongyan Zhao</i>	
MaFeRw: Query Rewriting with Multi-Aspect Feedbacks for Retrieval-Augmented Large Language Models	25434
<i>Yujing Wang, Hainan Zhang, Liang Pang, Binghui Guo, Hongwei Zheng, Zhiming Zheng</i>	
Defending Against Sophisticated Poisoning Attacks with RL-based Aggregation in Federated Learning	25443
<i>Yujing Wang, Hainan Zhang, Sijia Wen, Wangjie Qiu, Binghui Guo</i>	
RETQA: A Large-Scale Open-Domain Tabular Question Answering Dataset for Real Estate Sector	25452
<i>Zhensheng Wang, Wenmian Yang, Kun Zhou, Yiquan Zhang, Weijia Jia</i>	
STLC-KG:A Social Text Steganalysis Method Combining Large-Scale Language Models and Common-Sense Knowledge Graphs.....	25461
<i>Zhuang Wang, Linna Zhou, Xuekai Chen, Zhili Zhou, Zhongliang Yang</i>	
Is There No Such Thing as a Bad Question? H4R: HalluciBot for Ratiocination, Rewriting, Ranking, and Routing.....	25470
<i>William Watson, Nicole Cho, Nishan Srishankar</i>	
EBBS: An Ensemble with Bi-Level Beam Search for Zero-Shot Machine Translation.....	25479
<i>Yuqiao Wen, Behzad Shayegh, Chenyang Huang, Yanshuai Cao, Lili Mou</i>	
RLPF: Reinforcement Learning from Prediction Feedback for User Summarization with LLMs.....	25488
<i>Jiaxing Wu, Lin Ning, Luyang Liu, Harrison Lee, Neo Wu, Chao Wang, Sushant Prakash, Shawn O'Banion, Bradley Green, Jun Xie</i>	
TableBench: A Comprehensive and Complex Benchmark for Table Question Answering	25497
<i>Xianjie Wu, Jian Yang, Linzheng Chai, Ge Zhang, Jiaheng Liu, Xeron Du, Di Liang, Daixin Shu, Xianfu Cheng, Tianzhen Sun, Tongliang Li, Zhoujun Li, Guanglin Niu</i>	
Content-free Logical Modification of Large Language Model by Disentangling and Modifying Logic Representation.....	25507
<i>Xin Wu, Yuqi Bu, Yifei Chen, Yi Cai</i>	
Enhancing Audiovisual Speech Recognition Through Bifocal Preference Optimization.....	25516
<i>Yihan Wu, Yichen Lu, Yifan Peng, Xihua Wang, Ruihua Song, Shinji Watanabe</i>	
InverseCoder: Self-improving Instruction-Tuned Code LLMs with Inverse-Instruct	25525
<i>Yutong Wu, Di Huang, Wenxuan Shi, Wei Wang, Yewen Pu, Lingzhe Gao, Shihao Liu, Ziyuan Nan, Kaizhao Yuan, Rui Zhang, Xishan Zhang, Zidong Du, Qi Guo, Dawei Yin, Xing Hu, Yunji Chen</i>	
Improving Retrieval Augmented Language Model with Self-Reasoning.....	25534
<i>Yuan Xia, Jingbo Zhou, Zhenhui Shi, Jun Chen, Haifeng Huang</i>	
Detecting and Mitigating Hallucination in Large Vision Language Models via Fine-Grained AI Feedback.....	25543
<i>Wenyi Xiao, Ziwei Huang, Leilei Gan, Wanggui He, Haoyuan Li, Zhelun Yu, Fangxun Shu, Hao Jiang, Linchao Zhu</i>	

Explicitly Guided Difficulty-Controllable Visual Question Generation.....	25552
<i>Jiayuan Xie, Mengqiu Cheng, Xinting Zhang, Yi Cai, Guimin Hu, Mengying Xie, Qing Li</i>	
DSGram: Dynamic Weighting Sub-Metrics for Grammatical Error Correction in the Era of Large Language Models	25561
<i>Jinxiang Xie, Yilin Li, Xunjian Yin, Xiaojun Wan</i>	
Harnessing Large Language Models for Knowledge Graph Question Answering via Adaptive Multi-Aspect Retrieval-Augmentation	25570
<i>Derong Xu, Xinhang Li, Ziheng Zhang, Zhenxi Lin, Zhihong Zhu, Zhi Zheng, Xian Wu, Xiangyu Zhao, Tong Xu, Enhong Chen</i>	
Mitigating Social Bias in Large Language Models: A Multi-Objective Approach Within a Multi-Agent Framework.....	25579
<i>Zhenjie Xu, Wenqing Chen, Yi Tang, Xuanying Li, Cheng Hu, Zhixuan Chu, Kui Ren, Zibin Zheng, Zhichao Lu</i>	
S ³ Math: Spontaneous Step-Level Self-Correction Makes Large Language Models Better Mathematical Reasoners.....	25588
<i>Yuchen Yan, Jin Jiang, Yang Liu, Yixin Cao, Xin Xu, Mengdi Zhang, Xunliang Cai, Jian Shao</i>	
SongEditor: Adapting Zero-Shot Song Generation Language Model as a Multi-Task Editor.....	25597
<i>Chenyu Yang, Shuai Wang, Hangting Chen, Jianwei Yu, Wei Tan, Rongzhi Gu, Yaoxun Xu, Yizhi Zhou, Haina Zhu, Haizhou Li</i>	
Improving Factuality in Large Language Models via Decoding-Time Hallucinatory and Truthful Comparators	25606
<i>Dingkang Yang, Dongling Xiao, Jinjie Wei, Mingcheng Li, Zhaoyu Chen, Ke Li, Lihua Zhang</i>	
UniMuMo: Unified Text, Music, and Motion Generation.....	25615
<i>Han Yang, Kun Su, Yutong Zhang, Jiaben Chen, Kaizhi Qian, Gaowen Liu, Chuang Gan</i>	
Fine-Tuning Language Models with Collaborative and Semantic Experts	25624
<i>Jiaxi Yang, Binyuan Hui, Min Yang, Jian Yang, Lei Zhang, Qiang Qu, Junyang Lin</i>	
RAZOR: Sharpening Knowledge by Cutting Bias with Unsupervised Text Rewriting	25633
<i>Shuo Yang, Bardh Prenkaj, Gjergji Kasneci</i>	
MSE-Adapter: A Lightweight Plugin Endowing LLMs with the Capability to Perform Multimodal Sentiment Analysis and Emotion Recognition	25642
<i>Yang Yang, Xunde Dong, Yupeng Qiang</i>	
Is Sarcasm Detection a Step-by-Step Reasoning Process in Large Language Models?	25651
<i>Ben Yao, Yazhou Zhang, Qiuchi Li, Jing Qin</i>	
ReFF: Reinforcing Format Faithfulness in Language Models Across Varied Tasks	25660
<i>Jiashu Yao, Heyan Huang, Zeming Liu, Haoyu Wen, Wei Su, Boao Qian, Yuhang Guo</i>	
StableVC: Style Controllable Zero-Shot Voice Conversion with Conditional Flow Matching.....	25669
<i>Jixun Yao, Yang Yuguang, Yu Pan, Ziqian Ning, Jianhao Ye, Hongbin Zhou, Lei Xie</i>	
EXCGEC: A Benchmark for Edit-Wise Explainable Chinese Grammatical Error Correction.....	25678
<i>Jingheng Ye, Shang Qin, Yinghui Li, Xuxin Cheng, Libo Qin, Hai-Tao Zheng, Ying Shen, Peng Xing, Zishan Xu, Guo Cheng, Wenhao Jiang</i>	

PART 2

Assessing the Creativity of LLMs in Proposing Novel Solutions to Mathematical Problems	25687
<i>Junyi Ye, Jingyi Gu, Xinyun Zhao, Wenpeng Yin, Guiling Wang</i>	
Codec Does Matter: Exploring the Semantic Shortcoming of Codec for Audio Language Model	25697
<i>Zhen Ye, Peiwen Sun, Jiahe Lei, Hongzhan Lin, Xu Tan, Zheqi Dai, Qiuqiang Kong, Jianyi Chen, Jiahao Pan, Qifeng Liu, Yike Guo, Wei Xue</i>	
NLSR: Neuron-Level Safety Realignment of Large Language Models Against Harmful Fine-Tuning	25706
<i>Xin Yi, Shunfan Zheng, Linlin Wang, Gerard de Melo, Xiaoling Wang, Liang He</i>	
MDD-5k: A New Diagnostic Conversation Dataset for Mental Disorders Synthesized via Neuro-Symbolic LLM Agents	25715
<i>Congchi Yin, Feng Li, Shu Zhang, Zike Wang, Jun Shao, Piji Li, Jianhua Chen, Xun Jiang</i>	
MABR: Multilayer Adversarial Bias Removal Without Prior Bias Knowledge.....	25724
<i>Maxwell J. Yin, Boyu Wang, Charles Ling</i>	
Heuristic-free Knowledge Distillation for Streaming ASR via Multi-modal Training	25733
<i>Ji Won Yoon</i>	
SongGLM: Lyric-to-Melody Generation with 2D Alignment Encoding and Multi-Task Pre-Training	25742
<i>Jiaxing Yu, Xinda Wu, Yunfei Xu, Tiejiao Zhang, Songruoyao Wu, Le Ma, Kejun Zhang</i>	
Training with “Paraphrasing the Original Text” Teaches LLM to Better Retrieve in Long-Context Tasks.....	25751
<i>Yijiong Yu, Yongfeng Huang, Zhixiao Qi, Zhe Zhou</i>	
DMT-RoleBench: A Dynamic Multi-Turn Dialogue Based Benchmark for Role-Playing Evaluation of Large Language Model and Agent	25760
<i>Dingbo Yuan, Yipeng Chen, Guodong Liu, Chenchen Li, Chengfu Tang, Dongxu Zhang, Zhenkui Wang, Xudong Wang, Song Liu</i>	
Towards Robust Knowledge Unlearning: An Adversarial Framework for Assessing and Improving Unlearning Robustness in Large Language Models	25769
<i>Hongbang Yuan, Zhuoran Jin, Pengfei Cao, Yubo Chen, Kang Liu, Jun Zhao</i>	
CogSQL: A Cognitive Framework for Enhancing Large Language Models in Text-to-SQL Translation.....	25778
<i>Hongwei Yuan, Xiu Tang, Ke Chen, Lidan Shou, Gang Chen, Huan Li</i>	
Improving Natural Language Understanding for LLMs via Large-Scale Instruction Synthesis	25787
<i>Lin Yuan, Jun Xu, Honghao Gui, Mengshu Sun, Zhiqiang Zhang, Lei Liang, Jun Zhou</i>	
Synergistic Multi-Agent Framework with Trajectory Learning for Knowledge-Intensive Tasks.....	25796
<i>Shengbin Yue, Siyuan Wang, Wei Chen, Xuanjing Huang, Zhongyu Wei</i>	
QJL: 1-Bit Quantized JL Transform for KV Cache Quantization with Zero Overhead	25805
<i>Amir Zandieh, Majid Daliri, Insu Han</i>	
Self-Supervised Collaborative Information Bottleneck for Text Readability Assessment	25814
<i>Jinshan Zeng, Xianglong Yu, Xianchao Tong, Wenyan Xiao</i>	

Dual-View Learning for Conversational Emotion Recognition Through Context and Emotion-Shift Modeling	25823
<i>Xupeng Zha, Huan Zhao, Guanghui Ye, Zixing Zhang</i>	
Aligning Language Models Using Follow-up Likelihood as Reward Signal	25832
<i>Chen Zhang, Dading Chong, Feng Jiang, Chengguang Tang, Anningzhe Gao, Guohua Tang, Haizhou Li</i>	
Causal Prompting: Debiasing Large Language Model Prompting Based on Front-Door Adjustment	25842
<i>Congzhi Zhang, Linhai Zhang, Jialong Wu, Yulan He, Deyu Zhou</i>	
Controlling Large Language Models Through Concept Activation Vectors.....	25851
<i>Hanyu Zhang, Xiting Wang, Chengao Li, Xiang Ao, Qing He</i>	
CSR:Achieving 1 Bit Key-Value Cache via Sparse Representation.....	25860
<i>Hongxuan Zhang, Yao Zhao, Jiaqi Zheng, Chenyi Zhuang, Jinjie Gu, Guihai Chen</i>	
SAIL: Sample-Centric In-Context Learning for Document Information Extraction.....	25868
<i>Jinyu Zhang, Zhiyuan You, Jize Wang, Xinyi Le</i>	
Quantum-inspired Non-homologous Representation Constraint Mechanism for Long-tail Senses of Word Sense Disambiguation	25877
<i>Junwei Zhang, Xiaolin Li</i>	
Hierarchical Context Pruning: Optimizing Real-World Code Completion with Repository-Level Pretrained Code LLMs	25886
<i>Lei Zhang, Yunshui Li, Jiaming Li, Xiaobo Xia, Jiayi Yang, Run Luo, Minzheng Wang, Longze Chen, Junhao Liu, Qiang Qu, Min Yang</i>	
KnowPO: Knowledge-Aware Preference Optimization for Controllable Knowledge Selection in Retrieval-Augmented Language Models	25895
<i>Ruizhe Zhang, Yongxin Xu, Yuzhen Xiao, Runchuan Zhu, Xinke Jiang, Xu Chu, Junfeng Zhao, Yasha Wang</i>	
Look Before You Leap: Enhance Attention and Vigilance Regarding Harmful Content with GuidelineLLM.....	25904
<i>Shaoqing Zhang, Zhuosheng Zhang, Kehai Chen, Rongxiang Weng, Muyun Yang, Tiejun Zhao, Min Zhang</i>	
Code-switching Mediated Sentence-level Semantic Learning	25913
<i>Shuai Zhang, Jiangyan Yi, Zhengqi Wen, Jianhua Tao, Feihu Che, Jinyang Wu, Ruiibo Fu</i>	
FaceSpeak: Expressive and High-Quality Speech Synthesis from Human Portraits of Different Styles	25922
<i>Tian-Hao Zhang, Jiawei Zhang, Jun Wang, Xinyuan Qian, Xu-Cheng Yin</i>	
TrustUQA: A Trustful Framework for Unified Structured Data Question Answering	25931
<i>Wen Zhang, Long Jin, Yushan Zhu, Jiaoyan Chen, Zhiwei Huang, Junjie Wang, Yin Hua, Lei Liang, Huajun Chen</i>	
MEPNet: Medical Entity-Balanced Prompting Network for Brain CT Report Generation.....	25940
<i>Xiaodan Zhang, Yanzhao Shi, Junzhong Ji, Chengxin Zheng, Liangqiong Qu</i>	
Unveiling the Impact of Coding Data Instruction Fine-Tuning on Large Language Models Reasoning	25949
<i>Xinlu Zhang, Zhiyu Zoey Chen, Xi Ye, Xianjun Yang, Lichang Chen, William Yang Wang, Linda Ruth Petzold</i>	

Capability Instruction Tuning.....	25958
<i>Yi-Kai Zhang, De-Chuan Zhan, Han-Jia Ye</i>	
Multi-Attribute Multi-Grained Adaptation of Pre-Trained Language Models for Text Understanding from Bayesian Perspective	25967
<i>You Zhang, Jin Wang, Liang-Chih Yu, Dan Xu, Xuejie Zhang</i>	
Defeasible Visual Entailment: Benchmark, Evaluator, and Reward-Driven Optimization	25976
<i>Yue Zhang, Liqiang Jing, Vibhav Gogate</i>	
Enhancing Chain of Thought Prompting in Large Language Models via Reasoning Patterns	25985
<i>Yufeng Zhang, Xuepeng Wang, Lingxiang Wu, Jinqiao Wang</i>	
Encoder of Thoughts: Enhancing Planning Ability in Language Agents Through Structural Embedding.....	25994
<i>Yuxiang Zhang, Jitao Sang</i>	
Prototype-Guided Multimodal Relation Extraction based on Entity Attributes	26003
<i>Zefan Zhang, Weiqi Zhang, Yanhui Li, Tian Bai</i>	
Mixture of Knowledge Minigraph Agents for Literature Review Generation.....	26012
<i>Zhi Zhang, Yan Liu, Sheng-hua Zhong, Gong Chen, Yu Yang, Jiannong Cao</i>	
DiveR-CT: Diversity-enhanced Red Teaming Large Language Model Assistants with Relaxing Constraints.....	26021
<i>Andrew Zhao, Quentin Xu, Matthieu Lin, Shenchi Wang, Yong-Jin Liu, Zilong Zheng, Gao Huang</i>	
Beyond IID: Optimizing Instruction Finetuning from the Perspective of Instruction Interaction and Dependency	26031
<i>Hanyu Zhao, Li Du, Yiming Ju, Chengwei Wu, Tengfei Pan</i>	
CareBot: A Pioneering Full-Process Open-Source Medical Language Model	26039
<i>Lulu Zhao, Weihao Zeng, Xiaofeng Shi, Hua Zhou</i>	
Leveraging Attention to Effectively Compress Prompts for Long-Context LLMs	26048
<i>Yunlong Zhao, Haoran Wu, Bo Xu</i>	
Multi-Granular Multimodal Clue Fusion for Meme Understanding	26057
<i>Li Zheng, Hao Fei, Ting Dai, Zuquan Peng, Fei Li, Huisheng Ma, Chong Teng, Donghong Ji</i>	
MRR-FV: Unlocking Complex Fact Verification with Multi-Hop Retrieval and Reasoning	26066
<i>Liwen Zheng, Chaozhuo Li, Litian Zhang, Haoran Jia, Senzhang Wang, Zheng Liu, Xi Zhang</i>	
Hierarchical Divide-and-Conquer for Fine-Grained Alignment in LLM-Based Medical Evaluation	26075
<i>Shunfan Zheng, Xiechi Zhang, Gerard de Melo, Xiaoling Wang, Linlin Wang</i>	
BindGPT: A Scalable Framework for 3D Molecular Design via Language Modeling and Reinforcement Learning	26083
<i>Artem Zholus, Maksim Kuznetsov, Roman Schutski, Rim Shayakhmetov, Daniil Polykovskiy, Sarath Chandar, Alex Zhavoronkov</i>	
MoE-LPR: Multilingual Extension of Large Language Models Through Mixture-of-Experts with Language Priors Routing	26092
<i>Hao Zhou, Zhijun Wang, Shujian Huang, Xin Huang, Xue Han, Junlan Feng, Chao Deng, Weihua Luo, Jiajun Chen</i>	

CharacterBench: Benchmarking Character Customization of Large Language Models	26101
<i>Jinfeng Zhou, Yongkang Huang, Bosi Wen, Guanqun Bi, Yuxuan Chen, Pei Ke, Zhuang Chen, Xiyao Xiao, Libiao Peng, Kuntian Tang, Rongsheng Zhang, Le Zhang, Tangjie Lv, Zhipeng Hu, Hongning Wang, Minlie Huang</i>	
Editing Memories Through Few Targeted Neurons.....	26111
<i>Wei Zhou, Wei Wei, Guibang Cao, Fei Wang</i>	
Advancing Audio-Based Text Generation with Imbalance Preference Optimization.....	26120
<i>Zhenghao Zhou, Yongjie Liu, Chen Cao</i>	
RaDIO: Real-Time Hallucination Detection with Contextual Index Optimized Query Formulation for Dynamic Retrieval Augmented Generation	26129
<i>Jia Zhu, Hanghui Guo, Weijie Shi, Zhangze Chen, Pasquale De Meo</i>	
EventSum: A Large-Scale Event-Centric Summarization Dataset for Chinese Multi-News Documents.....	26138
<i>Mengna Zhu, Kaisheng Zeng, Mao Wang, Kaiming Xiao, Lei Hou, Hongbin Huang, Juanzi Li</i>	
DOMAINEVAL: An Auto-Constructed Benchmark for Multi-Domain Code Generation.....	26148
<i>Qiming Zhu, Jialun Cao, Yaojie Lu, Hongyu Lin, Xianpei Han, Le Sun, Shing-Chi Cheung</i>	
Utilize the Flow Before Stepping into the Same River Twice: Certainty Represented Knowledge Flow for Refusal-Aware Instruction Tuning	26157
<i>Runchuan Zhu, Zhipeng Ma, Jiang Wu, Junyuan Gao, Jiaqi Wang, Dahua Lin, Conghui He</i>	
One Token Can Help! Learning Scalable and Pluggable Virtual Tokens for Retrieval-Augmented Large Language Models	26166
<i>Yutao Zhu, Zhaocheng Huang, Zhicheng Dou, Ji-Rong Wen</i>	
PokerBench: Training Large Language Models to Become Professional Poker Players.....	26175
<i>Richard Zhuang, Akshat Gupta, Richard Yang, Aniket Rahane, Zhengyu Li, Gopala Anumanchipalli</i>	
Math-PUMA: Progressive Upward Multimodal Alignment to Enhance Mathematical Reasoning	26183
<i>Wenwen Zhuang, Xin Huang, Xiantao Zhang, Jin Zeng</i>	
TNCSE: Tensor Norm Constraints for Unsupervised Contrastive Learning of Sentence Embeddings.....	26192
<i>Tianyu Zong, Bingkan Shi, Hongzhu Yi, Jungang Xu</i>	

Author Index