

Haoran Luo

Created: 2026-04-12 06:06:59Z

(+65) 87626691, (+86) 15801159812
haoran.luo@ieee.org

Research Fellow, NTU Singapore
N4-B1a-02, Artificial Intelligence Lab, CCDS, NTU

ACADEMIC BACKGROUND

- **College of Computing and Data Science, Nanyang Technological University** Singapore, SG
Research Fellow, Advisor: Luu Anh Tuan Jan. 2026 – Present
- **College of Computing and Data Science, Nanyang Technological University** Singapore, SG
Joint Ph.D. in Artificial Intelligence, Advisor: Luu Anh Tuan Nov. 2024 – Oct. 2025
- **Computer Science, Beijing University of Posts and Telecommunications** Beijing, CN
Ph.D. in Computer Science & Technology, Advisor: Haihong E Sep. 2021 – Nov. 2025
- **Computer Science, Beijing University of Posts and Telecommunications** Beijing, CN
B.Eng. in Computer Science & Technology Sep. 2017 – Jun. 2021

SELECTED PUBLICATIONS

* Equal contribution. † Corresponding author.

Preprint

- Fangyu Hao, Jiayu Yang, Yifan Zhu[†], Zijun Yu, Qicen Wu, Wang Yunlong, Jiawei Li, Yulin Liu, Xu Zeng, Guanting Chen, Shihao Li, Zhonghong Ou, Meina Song, Mengyang Sun, **Haoran Luo**[†], Yu Shi, Yingyi Wang. *LungCURE: Benchmarking Multimodal Real-World Clinical Reasoning for Precision Lung Cancer Diagnosis and Treatment*. arXiv:2604.06925. [preprint] [code]
- Yandan Zheng, **Haoran Luo**[†], Zhenghong Lin, Wenjin Liu, Luu Anh Tuan[†]. *BenchBench: Benchmarking Automated Benchmark Generation*. arXiv:2603.20807. [preprint] [code]
- Ziyue Zhu, Shangyang Wu, Shuai Zhao[†], Zhiqiu Zhao, Shengjie Li, Yi Wang, Fang Li, **Haoran Luo**[†]. *NS-VLA: Towards Neuro-Symbolic Vision-Language-Action Models*. arXiv:2603.09542. [preprint] [code]
- Yifan Zhu, Guanting Chen, Bing Wei, **Haoran Luo**[†]. *HiFlow: Hierarchical Feedback-Driven Optimization for Constrained Long-Form Text Generation*. arXiv:2603.04996. [preprint]
- Yichao Feng, **Haoran Luo**[†], Zhenghong Lin, Yiqun Sun, Pengfei Wei, Lawrence B. Hsieh, Anh Tuan Luu[†]. *OrchMAS: Orchestrated Reasoning with Multi Collaborative Heterogeneous Scientific Expert Structured Agents*. arXiv:2603.03005. [preprint] [code]
- Yongqing Jiang, Jianze Wang, Zhiqi Shen, Zhenghong Lin, Jiayuan Wang, Yijian Yang, Kaoshan Dai[†], **Haoran Luo**[†]. *Rethinking Scientific Modeling: Toward Physically Consistent and Simulation-Executable Programmatic Generation*. arXiv:2602.07083. [preprint] [code]
- Fangzhi Xu, Hang Yan, Qiushi Sun, Jinyang Wu, Zixian Huang, Muye Huang, Jingyang Gong, Zichen Ding, Kanzhi Cheng, Yian Wang, Xinyu Che, Zeyi Sun, Jian Zhang, Zhangyue Yin, **Haoran Luo**, Xuanjing Huang, Ben Kao, Jun Liu, Qika Lin. *OdysseyArena: Benchmarking Large Language Models For Long-Horizon, Active and Inductive Interactions*. arXiv:2602.05843. [preprint] [code]
- Yifan Zhu, Xinyu Mu, Tao Feng, Zhonghong Ou, Yuning Gong, **Haoran Luo**[†]. *OmniRAG-Agent: Agentic Omnimodal Reasoning for Low-Resource Long Audio-Video Question Answering*. arXiv:2602.03707. [preprint]
- Mingda Zhang^{*}, **Haoran Luo**^{*}, Tiesunlong Shen, Qika Lin, Xiaoying Tang, Rui Mao, Erik Cambria. *FlowSteer: Interactive Agentic Workflow Orchestration via End-to-End Reinforcement Learning*. arXiv:2602.01664. [preprint] [code]
- Jinyang Wu, Changpeng Yang, Yuhao Shen, Fangzhi Xu, Bolin Ni, Chonghua Liao, Yuchen Liu, Hongzhen Wang, Shuai Nie, Shuai Zhang, **Haoran Luo**[†], Jiaming Xu[†]. *SSL: Sweet Spot Learning for Differentiated Guidance in Agentic Optimization*. arXiv:2601.22491. [preprint]
- Jiapu Wang, Xinghe Cheng, Zezheng Wu, Ruiqi Ma, Rui Wang, Zhichao Yan, **Haoran Luo**, Yuhao Jiang, Kai Sun. *Cumulative Path-Level Semantic Reasoning for Inductive Knowledge Graph Completion*. arXiv:2601.05629. [preprint]

- Hongyu Lin, Haolin Pan, **Haoran Luo**[†], Yuchen Li, Kaichun Yao, Libo Zhang, Mingjie Xing[†], Yanjun Wu. *AwareCompiler: Agentic Context-Aware Compiler Optimization via a Synergistic Knowledge-Data Driven Framework*. arXiv:2510.11759. [preprint] [code]
- Chengwei Wu, Jiapu Wang, Mingyang Gao, Xingrui Zhuo, Jipeng Guo, Runlin Lei, **Haoran Luo**, Tianyu Chen, Haoyi Zhou, Shirui Pan, Zechao Li. *CDTP: A Large-Scale Chinese Data-Text Pair Dataset for Comprehensive Evaluation of Chinese LLMs*. arXiv:2510.06039. [preprint]
- Jianheng Zhou, Chenyu Liu, Jinan Zhou, Yi Ding, Yang Liu, **Haoran Luo**, Ziyu Jia, Xinliang Zhou. *Introducing Multimodal Paradigm for Learning Sleep Staging PSG via General-Purpose Model*. arXiv:2509.22810. [preprint]
- Qika Lin, Yifan Zhu, Bin Pu, Ling Huang, **Haoran Luo**, Jingying Ma, Zhen Peng, Tianzhe Zhao, Fangzhi Xu, Jian Zhang, Kai He, Zhonghong Ou, Swapnil Mishra, Mengling Feng. *A Foundation Model for Chest X-ray Interpretation with Grounded Reasoning via Online Reinforcement Learning*. arXiv:2509.03906. [preprint] [code]
- Yi Wang^{*}, **Haoran Luo**^{*}, Lu Meng, Ziyu Jia, Xinliang Zhou, Qingsong Wen. *EEG-MedRAG: Enhancing EEG-based Clinical Decision-Making via Hierarchical Hypergraph Retrieval-Augmented Generation*. arXiv:2508.13735. [preprint] [code]
- Hongyu Lin, Yuchen Li, **Haoran Luo**[†], Kaichun Yao, Libo Zhang, Mingjie Xing[†], Yanjun Wu. *OS-R1: Agentic Operating System Kernel Tuning with Reinforcement Learning*. arXiv:2508.12551. [preprint] [code]
- **Haoran Luo**, Haihong E, Guanting Chen, Qika Lin, Yikai Guo, Fangzhi Xu, Zemin Kuang, Meina Song, Xiaobao Wu, Yifan Zhu, Luu Anh Tuan. *Graph-R1: Towards Agentic GraphRAG Framework via End-to-end Reinforcement Learning*. arXiv:2507.21892. [preprint] [code]
- Hongyu Lin, Yuchen Li, **Haoran Luo**[†], Kaichun Yao, Libo Zhang, Zhenghong Lin, Mingjie Xing[†], Yanjun Wu, Carl Yang. *BYOS: Knowledge-driven Large Language Models Bring Your Own Operating System More Excellent*. arXiv:2503.09663. [preprint] [code]
- Hongyu Lin^{*}, **Haoran Luo**^{*}, Hanghang Cao^{*}, Yang Liu, Shihao Gao, Kaichun Yao, Libo Zhang, Mingjie Xing, Yanjun Wu. *KGCompiler: Deep Learning Compilation Optimization for Knowledge Graph Complex Logical Query Answering*. arXiv:2503.02172. [preprint] [code]
- Shuai Lyu^{*}, **Haoran Luo**^{*}, Ripeng Li, Zhonghong Ou, Jiangfeng Sun, Yang Qin, Xiaoran Shang, Meina Song, Yifan Zhu. *SQL-o1: A Self-Reward Heuristic Dynamic Search Method for Text-to-SQL*. arXiv:2502.11741. [preprint] [code]

2026

1. Wenjin Liu, **Haoran Luo**[†], Xueyuan Lin, Haoming Liu[†], Tiesunlong Shen, Jiapu Wang, Rui Mao, Erik Cambria. *Prompt-R1: Collaborative Automatic Prompting Framework via End-to-end Reinforcement Learning*. The Annual Meeting of the Association for Computational Linguistics (**ACL**), 2026. [preprint] [code]
2. Wenjin Liu^{*}, **Haoran Luo**^{*}, Xin Feng, Xiang Ji, Lijuan Zhou, Rui Mao, Jiapu Wang, Shirui Pan, Erik Cambria. *LexGenius: An Expert-Level Benchmark for Large Language Models in Legal General Intelligence*. The Annual Meeting of the Association for Computational Linguistics (**ACL**), 2026. [preprint] [code]
3. Jinyang Wu, Mingkuan Feng, Shuai Zhang, Feihu Che, Zengqi Wen, Chonghua Liao, Ling Yang, **Haoran Luo**, Zheng Lian, Jianhua Tao. *Beyond Examples: Towards Automated Thought-level In-Context Reasoning for Large Language Models*. The Annual Meeting of the Association for Computational Linguistics (**ACL**), 2026. [preprint] [code]
4. Jinyang Wu, Chonghua Liao, Mingkuan Feng, Shuai Zhang, Zhengqi Wen, **Haoran Luo**, Ling Yang, Huazhe Xu, Jianhua Tao. *TemplateRL: Structured Template-Guided Reinforcement Learning for LLM Reasoning*. The Annual Meeting of the Association for Computational Linguistics (**ACL**), 2026. [preprint]
5. Hang Yan, Fangzhi Xu, Rongman Xu, Yifei Li, Jian Zhang, **Haoran Luo**, Xiaobao Wu, Luu Anh Tuan, Haiteng Zhao, Qika Lin, Jun Liu. *MUR: Momentum Uncertainty guided Reasoning for Large Language Models*. The Annual Meeting of the Association for Computational Linguistics (**ACL**), 2026. [preprint] [code]

6. Jian Zhang, Zhiyuan Wang, Zhangqi Wang, Yu He, **Haoran Luo**, Li Yuan, Lingling Zhang, Rui Mao, Qika Lin, Jun Liu. *MAXS: Meta-Adaptive Exploration with LLM Agents*. The Annual Meeting of the Association for Computational Linguistics (**ACL**), 2026. [preprint]
7. Hongyu Lin, Yang Liu, **Haoran Luo**, Linhai Li, Hanghan Cao, Mengna Wang, Jiahao Xu, Quanyi Li, Hongbin Zhang, Mingjie Xing, Yanjun Wu. *A Survey on Deep Learning Compilation Optimizations*. **Journal of Software**, 2026.
8. Yikai Guo, Xuemeng Tian, Bin Ge, Yuting Yang, Yao He, Wenjun Ke, Junjie Hu, Yanyang Li, **Haoran Luo**. *ConsistEAE: Enhancing Low-Resource Event Argument Extraction with Linguistically Consistent Demonstrations*. **Neurocomputing**, 2026. [paper]
9. Yifan Zhu, Huiqiang Rong, **Haoran Luo**[†]. *Token-Guard: Towards Token-Level Hallucination Control via Self-Checking Decoding*. The International Conference on Learning Representations (**ICLR**), 2026. [paper] [code]
10. Feng Yichao*, **Haoran Luo***, Lang Feng, Shuai Zhao, Anh Tuan Luu. *From Stimuli to Minds: Enhancing Psychological Reasoning in LLMs via Bilateral Reinforcement Learning*. The Annual AAAI Conference on Artificial Intelligence (**AAAI**), 2026. [paper] [code]

2025

11. Rui Pan*, **Haoran Luo***, Quan Yuan, Guiyang Luo, Jinglin Li, Tiesunlong Shen, Rui Mao, Erik Cambria. *Multitask Reinforcement Learning with Metadata-Guided Adaptive Routing*. **Information Fusion**, 2025. [paper]
12. **Haoran Luo**, Haihong E, Guanting Chen, Yandan Zheng, Xiaobao Wu, Yikai Guo, Qika Lin, Yu Feng, Zemin Kuang, Meina Song, Yifan Zhu, Luu Anh Tuan. *HyperGraphRAG: Retrieval-Augmented Generation via Hypergraph-Structured Knowledge Representation*. The Annual Conference on Neural Information Processing Systems (**NeurIPS**), 2025. [preprint] [code]
13. Haolin Pan, Hongyu Lin, **Haoran Luo**[†], Yang Liu, Kaichun Yao, Libo Zhang, Mingjie Xing[†], Yanjun Wu. *Compiler-R1: Towards Agentic Compiler Auto-tuning with Reinforcement Learning*. The Annual Conference on Neural Information Processing Systems (**NeurIPS**), 2025. [preprint] [code]
14. Yikai Guo, Bin Wang[†], Zhixiang Jiang, Shuo Tian, Kairui Wang, Wenjun Ke, **Haoran Luo**[†]. *Robust Detection of AI-Generated Text: Insights on Evolving LLMs and Adversarial Data*. China Conference on Knowledge Graph and Semantic Computing (**CCKS**), 2025.
15. Jun Zhang, Haihong E, Tianyi Hu, Yifan Zhu, Meina Song, **Haoran Luo**. *Complex Numerical Reasoning with Numerical Semantic Pre-training Framework*. The Conference on Empirical Methods in Natural Language Processing (**EMNLP**), 2025. [paper] [code]
16. Kaiyang Wan, Honglin Mu, Rui Hao, **Haoran Luo**, Tianle Gu, Xiuying Chen. *A Cognitive Writing Perspective for Constrained Long-Form Text Generation*. The Annual Meeting of the Association for Computational Linguistics (**ACL**), 2025. [paper] [code]
17. **Haoran Luo**, Haihong E, Yikai Guo, Qika Lin, Xiaobao Wu, Xinyu Mu, Wenhao Liu, Meina Song, Yifan Zhu, Luu Anh Tuan. *KBQA-o1: Agentic Knowledge Base Question Answering with Monte Carlo Tree Search*. The International Conference on Machine Learning (**ICML**), 2025. [paper] [code]
18. Yu Feng, Yangli-ao Geng, Yifan Zhu, Zongfu Han, Xie Yu, Kaiwen Xue, **Haoran Luo**, Mengyang Sun, Guangwei Zhang, Meina Song. *PM-MOE: Mixture of Experts on Private Model Parameters for Personalized Federated Learning*. The ACM Web Conference (**WWW**), 2025. [paper] [code]
19. Ningyuan Li, Haihong E, Tianyu Yao, Tianyi Hu, Yuhua Li, **Haoran Luo**, Meina Song, Yifan Zhu. *INFER: A Neural-symbolic Model For Extrapolation Reasoning on Temporal Knowledge Graph*. The International Conference on Learning Representations (**ICLR**), 2025. [paper] [code]
20. Shuai Lyu, Zijing Tian, Zhonghong Ou, Yifan Zhu, Xiao Zhang, Qiankun Ha, **Haoran Luo**, Meina Song. *TSVC: Tripartite Learning with Semantic Variation Consistency for Robust Image-Text Retrieval*. The Annual AAAI Conference on Artificial Intelligence (**AAAI**), 2025. [paper]

2024

21. **Haoran Luo**, Haihong E, Yuhao Yang, Tianyu Yao, Yikai Guo, Zichen Tang, Wentai Zhang, Shiyao Peng, Kaiyang Wan, Meina Song, Wei Lin, Yifan Zhu, Luu Anh Tuan. *Text2NKG: Fine-Grained N-ary Relation Extraction for N-ary relational Knowledge Graph Construction*. The Annual Conference on Neural Information Processing Systems (**NeurIPS**), 2024. [paper] [code]
22. Yu Feng, Zhen Tian, Yifan Zhu, Zongfu Han, **Haoran Luo**, Guangwei Zhang, Meina Song. *CP-Prompt: Composition-Based Cross-modal Prompting for Domain-Incremental Continual Learning*. The ACM International Conference on Multimedia (**ACM MM**), 2024. [paper] [code]
23. **Haoran Luo**, Haihong E, Zichen Tang, Shiyao Peng, Yikai Guo, Wentai Zhang, Chenghao Ma, Guanting Dong, Meina Song, Wei Lin, Yifan Zhu, Luu Anh Tuan. *ChatKBQA: A Generate-then-Retrieve Framework for Knowledge Base Question Answering with Fine-tuned Large Language Models*. The Annual Meeting of the Association for Computational Linguistics (**ACL**), 2024. [paper] [code]
24. Wenjun Ke, Yikai Guo, Qi Liu, Wanyi Chen, Peng Wang, **Haoran Luo**, Zhizhao Luo. *MDM: Meta diffusion model for hard-constrained text generation*. Knowledge-Based Systems (**KBS**), 2024. [paper] [code]
25. Wentai Zhang, Haihong E, **Haoran Luo**, Mingzhi Sun. *FulBM: Fast fully batch maintenance for landmark-based 3-hop cover labeling*. ACM Transactions on Knowledge Discovery from Data (**ACM TKDD**), 2024. [paper] [code]

2023 and before

26. Xueyuan Lin, Haihong E, Chengjin Xu, Gengxian Zhou, **Haoran Luo**, Tianyi Hu, Fenglong Su, Ningyuan Li, Mingzhi Sun. *TFLEX: Temporal Feature-Logic Embedding Framework for Complex Reasoning over Temporal Knowledge Graph*. The Annual Conference on Neural Information Processing Systems (**NeurIPS**), 2023. [paper] [code]
27. Ningyuan Li, Haihong E, Shi Li, Mingzhi Sun, Tianyu Yao, Meina Song, Yong Wang, **Haoran Luo**. *TR-Rules: Rule-based Model for Link Forecasting on Temporal Knowledge Graph Considering Temporal Redundancy*. The Conference on Empirical Methods in Natural Language Processing (**EMNLP**), 2023. [paper] [code]
28. **Haoran Luo**, Haihong E, Yuhao Yang, Yikai Guo, Mingzhi Sun, Tianyu Yao, Zichen Tang, Kaiyang Wan, Meina Song, Wei Lin. *HAHE: Hierarchical Attention for Hyper-Relational Knowledge Graphs in Global and Local Level*. The Annual Meeting of the Association for Computational Linguistics (**ACL**), 2023. [paper] [code]
29. Xiaoxuan Xie, Haihong E, Zemin Kuang, Ling Tan, Gengxian Zhou, **Haoran Luo**, Jundi Li, Meina Song. *Triple-view Hyper-relational Knowledge Graph for Hypertension*. Journal of Chinese Information Processing (**JCIP**), 2023. [paper]
30. **Haoran Luo**, Haihong E, Yuhao Yang, Gengxian Zhou, Yikai Guo, Tianyu Yao, Zichen Tang, Xueyuan Lin, Kaiyang Wan. *NQE: N-ary Query Embedding for Complex Query Answering over Hyper-Relational Knowledge Graphs*. The Annual AAAI Conference on Artificial Intelligence (**AAAI**), 2023. [paper] [code]
31. **Haoran Luo**, Haihong E, Ling Tan, Gengxian Zhou, Tianyu Yao, Kaiyang Wan. *DHGE: Dual-View Hyper-Relational Knowledge Graph Embedding for Link Prediction and Entity Typing*. The Annual AAAI Conference on Artificial Intelligence (**AAAI**), 2023. [paper] [code]
32. Gengxian Zhou, Haihong E, Zemin Kuang, Ling Tan, Xiaoxuan Xie, Jundi Li, **Haoran Luo**. *Clinical decision support system for hypertension medication based on knowledge graph*. Computer Methods and Programs in Biomedicine (**CMPB**), 2022. [paper]

ACADEMIC SERVICES

-
- **Conference Area Chair:** NeurIPS 2025, NeurIPS 2026, KDD 2025, KDD 2026
 - **Conference PC Member/Reviewer:** NeurIPS 2024, NeurIPS 2025, ICML 2025, ICML 2026, ICLR 2025, ICLR 2026, CVPR 2026, ACL 2023, ACL 2024, ACL 2025, ACL 2026, AAAI 2023, AAAI 2024, AAAI 2025, AAAI 2026, KDD 2025, KDD 2026, WWW 2025, ACM MM 2025, ACM MM 2026, IJCAI 2026, ECCV 2026, EMNLP 2022, EMNLP 2023, EMNLP 2024, EMNLP 2025, NAACL 2024, NAACL 2025, COLING 2025,

CIKM 2024, AISTATS 2025, AISTATS 2026, PAKDD 2026, NLPCC 2024, NLPCC 2025, EACL 2026, CCKS 2025, CCKS 2026, IEEE ICKG 2025, BESC 2026, ACM MMAsia 2025

- **Journal Reviewer:** IEEE Transactions on Knowledge and Data Engineering (IEEE TKDE), IEEE Transactions on Dependable and Secure Computing (IEEE TDSC), IEEE Transactions on Cybernetics (IEEE TCYB), IEEE Transactions on Learning Technologies (IEEE TLT), ACM Transactions on Computing for Healthcare (ACM HEALTH), Transactions on Machine Learning Research (TMLR), Information Processing & Management (IP&M), Pattern Recognition (PR), Neural Networks, Journal of Systems & Software (JSS), Knowledge-Based Systems (KBS), Expert Systems With Applications (ESWA), Neurocomputing, ISPRS Journal of Photogrammetry and Remote Sensing (P&RS), International Journal of Medical Informatics (IJMEDI), Geo-spatial Information Science (TGSI), Biomedical Signal Processing and Control (BSPC)