

Haoran Luo

Created: 2025-09-03 09:26:18Z

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

Ph.D. Candidate, BUPT & NTU
N4-B1a-02, Artificial Intelligence Lab, CCDS, NTU

ACADEMIC BACKGROUND

- **College of Computing and Data Science, Nanyang Technological University** Singapore, SG
Joint Ph.D. in Artificial Intelligence, Advisor: [Luu Anh Tuan](#) Nov. 2024 – present
- **Computer Science, Beijing University of Posts and Telecommunications** Beijing, CN
Ph.D. in Computer Science & Technology, Advisor: [Haihong E](#) Sep. 2021 – present
- **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

- Yi Wang*, **Haoran Luo***, Lu Meng. *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. [\[preprint\]](#) [\[code\]](#)
- Feng Yichao*, **Haoran Luo***, Lang Feng, Shuai Zhao, Anh Tuan Luu. *From Stimuli to Minds: Enhancing Psychological Reasoning in LLMs via Bilateral Reinforcement Learning*. arXiv:2508.02458. [\[preprint\]](#)
- **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. [\[preprint\]](#) [\[code\]](#)
- 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*. arXiv:2507.14958. [\[preprint\]](#) [\[code\]](#)
- 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*. arXiv:2506.15701. [\[preprint\]](#) [\[code\]](#)
- **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*. arXiv:2503.21322. [\[preprint\]](#) [\[code\]](#)
- Hongyu Lin, Yuchen Li, **Haoran Luo**[†], Kaichun Yao, Libo Zhang, Mingjie Xing[†], Yanjun Wu. *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. [\[preprint\]](#) [\[code\]](#)

2025

1. Jun Zhang, Haihong E, Tianyi Hu, Yifan Zhu, Meina Song, **Haoran Luo**. *Complex Numerical Reasoning with Numerical Semantic Pre-training Framework*. Conference on Empirical Methods in Natural Language Processing (**EMNLP**), 2025.

2. Kaiyang Wan, Honglin Mu, Rui Hao, **Haoran Luo**, Tianle Gu, Xiuying Chen. *A Cognitive Writing Perspective for Constrained Long-Form Text Generation*. Annual Meeting of the Association for Computational Linguistics (**ACL**), 2025. [\[paper\]](#) [\[code\]](#)
3. **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*. Proceedings of the International Conference on Machine Learning (**ICML**), 2025. [\[preprint\]](#) [\[code\]](#)
4. 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*. Proceedings of the ACM on Web Conference (**WWW**), 2025. [\[paper\]](#) [\[code\]](#)
5. Ningyuan Li, Haihong E, Tianyu Yao, Tianyi Hu, Yuhao Li, **Haoran Luo**, Meina Song, Yifan Zhu. *INFER: A Neural-symbolic Model For Extrapolation Reasoning on Temporal Knowledge Graph*. Proceedings of the International Conference on Learning Representations (**ICLR**), 2025. [\[paper\]](#) [\[code\]](#)
6. 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*. Proceedings of the AAAI Conference on Artificial Intelligence (**AAAI**), 2025. [\[paper\]](#)

2024

7. **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*. Proceedings of the Advances in Neural Information Processing Systems (**NeurIPS**), 2024. [\[paper\]](#) [\[code\]](#)
8. 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*. Proceedings of the ACM International Conference on Multimedia (**ACM MM**), 2024. [\[paper\]](#) [\[code\]](#)
9. **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*. Annual Meeting of the Association for Computational Linguistics (**ACL**), 2024. [\[paper\]](#) [\[code\]](#)
10. 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. [\[paper\]](#) [\[code\]](#)
11. 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

12. 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*. Proceedings of the Advances in Neural Information Processing Systems (**NeurIPS**), 2023. [\[paper\]](#) [\[code\]](#)
13. 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*. Conference on Empirical Methods in Natural Language Processing (**EMNLP**), 2023. [\[paper\]](#) [\[code\]](#)
14. **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*. Annual Meeting of the Association for Computational Linguistics (**ACL**), 2023. [\[paper\]](#) [\[code\]](#)
15. **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*. Proceedings of the AAAI Conference on Artificial Intelligence (**AAAI**), 2023. [\[paper\]](#) [\[code\]](#)

16. **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*. Proceedings of the AAAI Conference on Artificial Intelligence (**AAAI**), 2023. [\[paper\]](#) [\[code\]](#)
17. 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**, 2022. [\[paper\]](#)

ACADEMIC SERVICES

- **Conference Area Chair:** NeurIPS 2025, KDD 2025, KDD 2026
- **Conference PC Member/Reviewer:** NeurIPS 2024, NeurIPS 2025, ICLR 2025, ICLR 2026, ICML 2025, ACL 2023, ACL 2024, ACL 2025, AAAI 2023, AAAI 2024, AAAI 2025, AAAI 2026, KDD 2025, WWW 2025, ACM MM 2025, EMNLP 2022, EMNLP 2023, EMNLP 2024, EMNLP 2025, NAACL 2024, NAACL 2025, COLING 2025, CIKM 2024, AISTATS 2025, AISTATS 2026, NLPCC 2024, NLPCC 2025, CCKS 2025, IEEE ICKG 2025, ACM MMAsia 2025
- **Journal Reviewer:** IEEE Transactions on Knowledge and Data Engineering (IEEE TKDE), IEEE Transactions on Dependable and Secure Computing (IEEE TDSC), Information Processing & Management (IPM), Pattern Recognition (PR), Knowledge-Based Systems (KBS), Neurocomputing, Geo-spatial Information Science (TGSI), Biomedical Signal Processing and Control (BSPC)