

QIANYUE HAO

✉ haoqianyue@hotmail.com · 🌐 <https://haohaoqian.github.io>

🎓 EDUCATION

Department of Electronic Engineering, Tsinghua University

2018.08 – 2022.06

Bachelor in Engineering, GPA 3.80/4.00 (Ranking Top 10%), Outstanding Thesis

Department of Electronic Engineering, Tsinghua University

2022.08 – Now

PhD in Engineering

📊 RESEARCH

Research Interests: Reinforcement Learning (RL), Large Language Models (LLMs)

- Reinforcement Learning in the post-training and inference-time of Large Language Models
Featured works: ICLR'26*2
- Large Language Models for textual data mining
Featured works: Nature, NeurIPS'25 (spotlight), NeurIPS'24
- Reinforcement Learning for large-scale applications
Featured works: KDD'25, KDD'24, KDD'23, KDD'22, KDD'21

Conference Publications

- [ICLR'26] RL of Thoughts: Navigating LLM Reasoning with Inference-time Reinforcement Learning
*Qianyue Hao**, *Sibo Li**, *Jian Yuan*, *Yong Li*
- [ICLR'26] Reinforcement Learning Fine-Tuning Enhances Activation Intensity and Diversity in the Internal Circuitry of LLMs
*Honglin Zhang**, *Qianyue Hao**, *Fengli Xu*, *Yong Li*
- [NeurIPS'25 (spotlight)] LLM-Explorer: A Plug-in Reinforcement Learning Policy Exploration Enhancement Driven by Large Language Models
*Qianyue Hao**, *Yiwen Song**, *Qingmin Liao*, *Jian Yuan*, *Yong Li*
- [NeurIPS'24] HLM-Cite: Hybrid Language Model Workflow for Text-based Scientific Citation Prediction
Qianyue Hao, *Jingyang Fan*, *Fengli Xu*, *Jian Yuan*, *Yong Li*
- [KDD'25] CoopRide: Cooperate All Grids in City-Scale Ride-Hailing Dispatching with Multi-Agent Reinforcement Learning
*Jingwei Wang**, *Qianyue Hao**, *Wenzhen Huang*, *Xiaochen Fan*, *Qin Zhang*, *Zhentao Tang*, *Bin Wang*, *Jianye Hao*, *Yong Li*
- [KDD'24] DyPS: Dynamic Parameter Sharing in Multi-Agent Reinforcement Learning for Spatio-Temporal Resource Allocation
*Jingwei Wang**, *Qianyue Hao**, *Wenzhen Huang*, *Xiaochen Fan*, *Zhentao Tang*, *Bin Wang*, *Jianye Hao*, *Yong Li*
- [KDD'23] GAT-MF: Graph Attention Mean Field for Very Large Scale Multi-Agent Reinforcement Learning
Qianyue Hao, *Wenzhen Huang*, *Tao Feng*, *Jian Yuan*, *Yong Li*
- [KDD'22] Reinforcement Learning Enhances the Experts: Large-scale COVID-19 Vaccine Allocation with Multi-factor Contact Network
Qianyue Hao, *Wenzhen Huang*, *Fengli Xu*, *Kun Tang*, *Yong Li*
- [KDD'21] Hierarchical Reinforcement Learning for Scarce Medical Resource Allocation with Imperfect Information
Qianyue Hao, *Fengli Xu*, *Lin Chen*, *Pan Hui*, *Yong Li*
- [KDD'20] Understanding the Urban Pandemic Spreading of COVID-19 with Real World Mobility Data
*Qianyue Hao**, *Lin Chen**, *Fengli Xu*, *Yong Li*
- [CIKM'25] CityLight: A Neighborhood-inclusive Universal Model for Coordinated City-scale Traffic Signal Control
Jinwei Zeng, *Chao Yu*, *Xinyi Yang*, *Wenxuan Ao*, *Qianyue Hao*, *Jian Yuan*, *Yong Li*, *Yu Wang*, *Huazhong Yang*

Journal Publications

- [Nature] Artificial Intelligence Tools Expand Scientists' Impact but Contract Science's Focus
Qianyue Hao, Fengli Xu, Yong Li, James Evans
- [Patterns] Toward Large Reasoning Models: A Survey of Reinforced Reasoning with Large Language Models
Fengli Xu, Qianyue Hao*, Chenyang Shao*, Zefang Zong*, Yu Li*, Jingwei Wang, Yunke Zhang, Jingyi Wang, Xiaocong Lan, Jiahui Gong, Tianjian Ouyang, Fanjin Meng, Yuwei Yan, Qinglong Yang, Yiwen Song, Sijian Ren, Xinyuan Hu, Jie Feng, Chen Gao, Yong Li*
- [TIST] Hierarchical Multi-agent Model for Reinforced Medical Resource Allocation with Imperfect Information
Qianyue Hao, Fengli Xu, Lin Chen, Pan Hui, Yong Li
- [CSUR] A Survey of Machine Learning for Urban Decision Making: Applications in Planning, Transportation, and Healthcare
Yu Zheng, Qianyue Hao, Jingwei Wang, Changzheng Gao, Jinwei Chen, Depeng Jin, Yong Li

Preprints

- [arxiv] KeyWorld: Key Frame Reasoning Enables Effective and Efficient World Models
Sibo Li, Qianyue Hao*, Yu Shang, Yong Li*
- [arxiv] Multiple Weaks Win Single Strong: Large Language Models Ensemble Weak Reinforcement Learning Agents into a Supreme One
Yiwen Song, Qianyue Hao*, Qingmin Liao, Jian Yuan, Yong Li*
- [ResearchGate] Reinforcement Learning in the Era of Large Language Models: Challenges and Opportunities
Qianyue Hao, Lin Chen, Xiaoqian Qi, Yuan Yuan, Zefang Zong, Hongyi Chen, Keyu Zhao, Shengyuan Wang, Yunke Zhang, Jian Yuan, Yong Li
- [arxiv] A Survey on Human-Centric LLMs
Jing Yi Wang, Nicholas Sukiennik, Tong Li, Weikang Su, Qianyue Hao, Jingbo Xu, Zihan Huang, Fengli Xu, Yong Li

🔗 INTERNSHIP

Schlumberger, Intern AI Engineer

2024.07 – 2024.08

- Geo-steering project that includes generation of training environments, RL model training and evaluation.
- High evaluation received from the project leader. US patent in application.

ByteDance, Intern AI Engineer

2021.06 – 2021.09

- Game AI project that includes RL model training and evaluation, as well as performance prediction of the agent.
- High evaluation received from the project leader.

🏆 AWARDS

Scholarship for Graduate Students, Tsinghua University	2025.10
Scholarship for Graduate Students, Department of Electronic Engineering, Tsinghua University	2023.10
Outstanding Undergraduate Thesis Award, Tsinghua University	2022.06
National Scholarship for Undergraduate Students, Ministry of Education of the PRC	2021.10
Scholarship for Undergraduate Students, Tsinghua University	2020.10