

Hao Hu

30 Shuangqing Road, MMW Building S219, Beijing, China

✉ huh22@mails.tsinghua.edu.cn 🌐 [mousehu.github.io/](https://github.com/mousehu)
📄 <https://scholar.google.com/citations?user=mhDH3VYAAAAJ>

Education

Tsinghua University Sept. 2019 - Present
Ph.D. Candidate in Computer Science, IIIS
Advisor: Prof. Chongjie Zhang and Prof. Yang Gao

Peking University Sept. 2015 - July 2019
B.S. in Theoretical and Applied Mechanics, School of Engineering
Double Major in Computer Science

Research Experience

NURL Group, IEMS, Northwestern University Apr. 2023 - Sept. 2023
Worked on reinforcement learning theory and reinforcement learning with large language models
Advisor: Prof. Zhaoran Wang

Machine Intelligence Group, IIIS, Tsinghua University Sept. 2019 - Present
Working on episodic memory, offline reinforcement learning, semi-supervised reinforcement learning
Advisor: Prof. Chongjie Zhang and Prof. Yang Gao

Vision and Media Computing Group, Peking University Nov. 2017 - Jun. 2018
Worked on semantic segmentation and autonomous driving
Advisor: Prof. Shiliang Zhang

Publications And Pre-Prints

(* indicates equal contribution)

1. Bayesian Design Principles for Offline-to-Online Reinforcement Learning

Hao Hu*, Yiqin Yang*, Jianing Ye, Chengjie Wu, Ziqing Mai, Yujing Hu, Tangjie Lv, Changjie Fan, Qianchuan Zhao, Chongjie Zhang

Forty-first International Conference on Machine Learning (ICML), 2024

2. Reason for Future, Act for Now: A Principled Architecture for Autonomous LLM Agents [Paper]

[Code] [Blog]

Zhihan Liu*, Hao Hu*, Shenao Zhang*, Hongyi Guo, Shuqi Ke, Boyi Liu, Zhaoran Wang

Forty-first International Conference on Machine Learning (ICML), 2024

NeurIPS Workshop on Foundation Models for Decision Making, 2023

3. Planning, Fast and Slow: Online Reinforcement Learning with Action-Free Offline Data via Multiscale Planners

Chengjie Wu*, Hao Hu*, Yiqin Yang, Ning Zhang, Chongjie Zhang

Forty-first International Conference on Machine Learning (ICML), 2024

4. Stylized Offline Reinforcement Learning: Extracting Diverse High-Quality Behaviors from Heterogeneous Datasets [Paper]

Yihuan Mao, Chengjie Wu, Xi Chen, Hao Hu, Ji Jiang, Tianze Zhou, Tangjie Lv, Changjie Fan, Zhipeng Hu, Yi Wu, Yujing Hu, Chongjie Zhang

Twelveth International Conference on Learning Representations (ICLR), 2024

5. **Unsupervised Behavior Extraction via Random Intent Priors** [\[Paper\]](#) [\[Code\]](#)

Hao Hu*, Yiqin Yang*, Jianing Ye, Ziqing Mai, Chongjie Zhang

Thirty-seventh Conference on Neural Information Processing Systems (NeurIPS), 2023

6. **One Objective to Rule Them All: A Maximization Objective Fusing Estimation and Planning for Exploration** [\[Paper\]](#) [\[Code\]](#)

Zhihan Liu*, Miao Lu*, Wei Xiong*, Han Zhong, Hao Hu, Shenao Zhang, Sirui Zheng, Zhuoran Yang, Zhaoran Wang

Thirty-seventh Conference on Neural Information Processing Systems (NeurIPS) Spotlight, 2023

7. **What is Essential for Unseen Goal Generalization of Offline Goal-conditioned RL?** [\[Paper\]](#) [\[Code\]](#)

Ruiyang, Yong Lin, Xiaoteng Ma, Hao Hu, Chongjie Zhang, Tong Zhang

Fortieth International Conference on Machine Learning (ICML), 2023

8. **The Provable Benefit of Unsupervised Data Sharing for Offline Reinforcement Learning** [\[Paper\]](#) [\[Code\]](#)

Hao Hu*, Yiqin Yang*, Qianchuan Zhao, Chongjie Zhang

Eleventh International Conference on Learning Representations (ICLR), 2023

9. **Flow to Control: Offline Reinforcement Learning with Lossless Primitive Discovery** [\[Paper\]](#) [\[Code\]](#)

Yiqin Yang*, Hao Hu*, Xiaoteng Ma*, Wenzhe Li*, Siyuan Li, Chongjie Zhang, Qianchuan Zhao

Thirty-Seventh AAAI Conference on Artificial Intelligence (AAAI), 2023

10. **On the Role of Discount Factor in Offline Reinforcement Learning** [\[Paper\]](#) [\[Code\]](#)

Hao Hu*, Yiqin Yang*, Qianchuan Zhao, Chongjie Zhang

Thirty-ninth International Conference on Machine Learning (ICML), 2022

11. **Offline Reinforcement Learning with Value-based Episodic Memory** [\[Paper\]](#) [\[Code\]](#)

Xiaoteng Ma*, Yiqin Yang*, Hao Hu*, Qihan Liu, Jun Yang, Chongjie Zhang, Qianchuan Zhao, Bin Liang

Tenth International Conference on Learning Representations (ICLR), 2022

12. **On the Estimation Bias in Double Q-Learning** [\[Paper\]](#) [\[Code\]](#)

Zhizhou Ren, Guangxiang Zhu, Hao Hu, Beining Han, Jianglun Chen, Chongjie Zhang

Thirty-fifth Conference on Neural Information Processing Systems (NeurIPS), 2021

13. **MetaCURE: Meta Reinforcement Learning with Empowerment-Driven Exploration** [\[Paper\]](#) [\[Code\]](#)

Jin Zhang*, Jianhao Wang*, Hao Hu, Tong Chen, Yingfeng Chen, Changjie Fan, Chongjie Zhang

Thirty-eighth International Conference on Machine Learning (ICML), 2021

14. **Generalizable Episodic Memory for Deep Reinforcement Learning** [\[Paper\]](#) [\[Code\]](#)

Hao Hu, Jianing Ye, Zhizhou Ren, Guangxiang Zhu, Chongjie Zhang

Thirty-eighth International Conference on Machine Learning (ICML), 2021

15. Query-Efficient Offline Preference-Based Reinforcement Learning via In-Dataset Exploration [\[Paper\]](#)

Hao Hu*, Yiqin Yang*, Shuai Wang, Bo Liu, Yang Gao, Chongjie Zhang

Under Review

Honors and Scholarships

Lingjun Pilot Scholarship	2023-2024
Huiyan Elite Scholarship	2022-2023
TOYOTA Scholarship	2021-2022
Second Prize in Zhou Peiyuan Mechanics Competition	MAY 2017
National Scholarship	2016-2017
Grand Prize in National College Students Physics Competition	SEPT. 2016
Award for Academic Excellents	2015-2016

Experience and Services

Engineering Intern

Microsoft STCA, News & Relevance team
Worked on newsletter recommendation

Jun. 2018 - Sept. 2018

Teaching Assistant

Deep Reinforcement Learning, Spring, 2022
Artificial Intelligence: Principles and Techniques, Fall, 2021
Artificial Intelligence: Principles and Techniques, Fall, 2020

Feb. 2022 - Jun. 2022
Sept. 2021 - Jan. 2022
Sept. 2020 - Jan. 2021

Student Instructor

Theoretical Mechanics, Spring, 2018

Feb. 2018 - Jun. 2018

Reviewer

International Conference on Machine Learning (ICML)
International Conference on Learning Representations (ICLR)
Conference on Neural Information Processing Systems (NeurIPS)
International Conference on Autonomous Agents and Multiagent Systems (AAMAS)

Selected Talks

- Data-Driven Reinforcement Learning [\[Slides\]](#) *Nov. 2023*
Bytedance AI Lab
- Unsupervised Behavior Extraction via Random Intent Priors [\[Slides\]](#) *Oct. 2023*
The 64th Seminar of Reinforcement Learning China Community (RL China)
- On the Role of Discount Factor in Offline Reinforcement Learning [\[Slides\]](#) [\[Video\]](#) *Jun. 2022*
The 21st Seminar of Reinforcement Learning China Community (RL China)
- Generalizable Episodic Memory for Deep Reinforcement Learning [\[Slides\]](#) *Jul. 2021*
Nanjing University