

胡昊



手机: (+86) 18811576155

邮箱: huh22@mails.tsinghua.edu.cn

个人主页: <https://mousehu.github.io/>

教育背景 EDUCATION

清华大学, 交叉信息研究院, 博士研究生在读
专业 计算机科学与技术

2019年9月 - 至今

北京大学, 工学院, 理学学士
专业 理论与应用力学

2015年9月 - 2019年7月

北京大学, 信息科学与技术学院, 理学学士
专业 计算机科学与技术 (双学位)

2016年9月 - 2019年7月

科研经历 RESEARCH

Machine Intelligence Group, 清华大学, 交叉信息研究院
导师 张崇洁 副教授 高阳 助理教授

2019年9月 - 至今

研究方向 深度强化学习, 强化学习理论, 离线强化学习

NURL Group, Northwestern University, IEMS
导师 汪昭然 助理教授

2023年4月 - 2023年9月

研究方向 强化学习理论

VMC Group, 北京大学, 信息科学与技术学院
导师 张史梁 副教授

2017年11月 - 2018年6月

研究方向 语义分割, 自动驾驶

实习经历 INTERNSHIP

微软亚洲工程院, 软件开发实习生

2018年6月 - 2018年9月

- 基于 Bing 点击数据的新闻文本分类与时事通讯推荐
- Sagitt 等内部工具开发

发表论文 PUBLICATIONS

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

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**
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**
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**
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, 2023)
7. **What is Essential for Unseen Goal Generalization of Offline Goal-conditioned RL?**
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**
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**
Yiqin Yang, Hao Hu*, 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**
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**
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**
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**
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**
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**

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

Under Review

获得荣誉 HONORS

灵均领航奖学金	2023-2024
惠妍英才奖学金	2022-2023
丰田奖学金 (一等)	2021-2022
北京大学优秀毕业生	2019.7
全国周培源力学竞赛二等奖	2017.5
国家奖学金	2016-2017
全国大学生物理竞赛 (部分赛区) 特等奖	2016.9
北京大学学习优秀奖	2015-2016

学术讲座 TALKS

- Data-Driven Reinforcement Learning 2023 年 11 月
字节跳动 AI Lab
- Unsupervised Behavior Extraction via Random Intent Priors 2023 年 10 月
The 64th Seminar of Reinforcement Learning China Community (RL China)
- On the Role of Discount Factor in Offline Reinforcement Learning 2022 年 6 月
The 21st Seminar of Reinforcement Learning China Community (RL China)
- Generalizable Episodic Memory for Deep Reinforcement Learning 2021 年 7 月
南京大学

教研服务 SERVICES

助教

深度强化学习, 2022 年春季学期	2022 年 2 月 - 2022 年 6 月
人工智能: 理论与方法, 2021 年秋季学期	2021 年 9 月 - 2022 年 1 月
人工智能: 理论与方法, 2020 年秋季学期	2020 年 9 月 - 2021 年 1 月

学生导师

理论力学, 2018 年春季学期	2018 年 2 月 - 2018 年 6 月
------------------	-------------------------

审稿人

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)