Jiayi Huang

Research Interests

My research interests lie at the intersection of statistics and machine learning, with a particular focus on reinforcement learning theory and applications. My work is dedicated to utilizing mathematical and statistical tools to establish the theoretical foundations of machine learning, aiming to guide the development of innovative algorithms for tackling complex real-world challenges.

EDUCATION

 University of Illinois Urbana-Champaign (UIUC) Siebel School of Computing and Data Science Ph.D. in <i>Computer Science</i> 	Aug. 2024 – Present Urbana, IL
Advisor: Prof. Nan Jiang	
Peking University (PKU) Center for Data Science	Sep. 2021 – Jul. 2024 Beijing, China
 M.S. in <i>Data Science (Statistics)</i> Advisor: Prof. Liwei Wang & Prof. Mohan Chen 	
University of Science and Technology of China (USTC) School of the Gifted Young (SGY)	Sep. 2017 – Jul. 2021 Hefei, China
 B.S. in Statistics Overall GPA: 3.99/4.3 (91.95) Rank: 2/75 in Statistics B.E. in Computer Science (Dual) Overall GPA: 3.90/4.3 (91.24) 	
 University of Washington (UW) Department of Electrical Engineering Summer School of <i>Global Electrical Engineering Program</i> 	Jul. 2018 – Aug. 2018 Seattle, WA

EXPERIENCE

University of California, Los Angeles (UCLA)

Research Intern, advised by Prof. Lin F. Yang

Mar. 2023 – Dec. 2023 *Remote*

- $\cdot\,$ Worked on reinforcement learning with heavy-tailed rewards.
 - \cdot We proposed two computationally efficient algorithms for heavy-tailed linear bandits and linear MDPs, based on a novel concentration inequality for adaptive Huber regression.
 - $\cdot\,$ These algorithms achieve both minimax optimal and instance-dependent regret bounds.
 - $\cdot\,$ We provided a lower bound to demonstrate the optimality.
 - $\cdot\,$ We also conducted numerical experiments to corroborate the computational efficiency.
- $\cdot\,$ Worked on reinforcement learning with general function approximation.

- \cdot We proposed an algorithm for model-based reinforcement learning with general function approximation, which features the novel combination of weighted value-targeted regression and a high-order moment estimator.
- · Our proposed algorithm achieves a both horizon-free and instance-dependent regret bound.
- It is both statistically and computationally efficient.
- We also conducted numerical experiments to validate the theoretical findings.

Peking University

Teaching Assistant	Beijing, China
Machine Learning (Turing Class)	Spring 2022
University of California, Los Angeles	Apr. 2021 – Sep. 2021

Research Intern, advised by Prof. Lin F. Yang

· Worked on linear bandits with super heavy-tailed rewards.

- · We proposed a generic algorithmic framework for super heavy-tailed linear bandits, which adopts a novel mean-of-medians estimator to handle the challenge of heavy-tailedness.
- We showed that our algorithmic framework is provably efficient for regret minimization.
- We also conducted numerical experiments to validate the effectiveness of our framework.

HONORS AND AWARDS

Peking University	
\cdot Graduate Special Academic Scholarship	Oct. 2022
University of Science and Technology of China	
\cdot Outstanding Graduate (USTC and Anhui Province)	Jun. 2021
\cdot National Scholarship	Dec. 2020
\cdot Nomination for Guo Moruo Scholarship	Oct. 2020
\cdot Outstanding Student Scholarship — Silver Award	Dec. 2018, 2019

PROFESSIONAL SERVICES

Conference Reviewer	AISTATS
Journal Reviewer	Artificial Intelligence, Machine Intelligence Research

PUBLICATIONS

- [1] Jiayi Huang, Han Zhong, Liwei Wang, Lin F. Yang. "Horizon-Free and Instance-Dependent Regret Bounds for Reinforcement Learning with General Function Approximation." AIS-TATS 2024. [Arxiv] [Code]
- [2] Jiavi Huang, Han Zhong, Liwei Wang, Lin F. Yang. "Tackling Heavy-Tailed Rewards in Reinforcement Learning with Function Approximation: Minimax Optimal and Instance-Dependent Regret Bounds." NeurIPS 2023. [Arxiv] [Code]
- [3] Han Zhong, Jiavi Huang, Lin F. Yang, Liwei Wang. "Breaking the Moments Condition Barrier: No-Regret Algorithm for Bandits with Super Heavy-Tailed Payoffs." NeurIPS 2021. [Arxiv] [Code]

Remote