

Zhihan Xiong

✉ zhihanx@cs.washington.edu

🏠 <https://homes.cs.washington.edu/~zhihanx/>

Education

University of Washington, Seattle, WA

Sep 2020–present

Ph.D. in Computer Science & Engineering, advisor: **Maryam Fazel**

GPA: 4.0/4.0

Stanford University, Stanford, CA

Sep 2018–Jun 2020

M.S. in Statistics

GPA: 4.0/4.0

University of Illinois at Urbana-Champaign, Champaign, IL

Aug 2014–May 2018

B.S. in Mathematics (Summa Cum Laude and with Highest Distinction)

GPA: 4.0/4.0

B.S. in Engineering Physics (with Highest Honors)

Minors in Computer Science and Statistics

Research Interests: Reinforcement Learning, Bandit Problems, Reinforcement Learning from Human Feedback (RLHF) in Large Language Models (LLMs).

Professional Services: Reviewer of ICML (2021, 2022, 2023, 2024), NeurIPS (2021, 2022, 2023) and ICLR (2022, 2023, 2024).

Publications/ Preprints ([Google Scholar](#))

(* indicates equal contributions)

Policy Mirror Descent with Dual Function Approximation

○ **Zhihan Xiong**, Maryam Fazel, Lin Xiao

○ *Preprint.*

A/B Testing and Best-arm Identification for Linear Bandits with Robustness to Non-stationarity

○ **Zhihan Xiong***, Romain Camilleri*, Maryam Fazel, Lalit Jain, Kevin Jamieson

○ *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2024. [[paper](#)]

○ *Conference on Digital Experimentation @ MIT (CODE@MIT)*, 2023

A Black-box Approach for Non-stationary Multi-agent Reinforcement Learning

○ Haozhe Jiang, Qiwen Cui, **Zhihan Xiong**, Maryam Fazel, Simon S. Du

○ *International Conference on Learning Representations (ICLR)*, 2024. [[arXiv](#)]

Offline Congestion Games: How Feedback Type Affects Data Coverage Requirement

○ Haozhe Jiang*, Qiwen Cui*, **Zhihan Xiong**, Maryam Fazel, Simon S. Du

○ *International Conference on Learning Representations (ICLR)*, 2023. [[paper](#)]

Learning in Congestion Games with Bandit Feedback

○ Qiwen Cui*, **Zhihan Xiong***, Maryam Fazel, Simon S. Du

○ *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. [[paper](#)]

Near-Optimal Randomized Exploration for Tabular Markov Decision Processes

○ **Zhihan Xiong***, Ruoqi Shen*, Qiwen, Cui*, Maryam Fazel, Simon S. Du

○ *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. [[paper](#)]

Fourier Learning with Cyclical Data

○ Yingxiang Yang*, **Zhihan Xiong***, Tianyi Liu*, Taiqing Wang, Chong Wang

- *International Conference on Machine Learning (ICML)*, 2022. [[paper](#)]

Selective Sampling for Online Best-arm Identification

- Romain Camilleri*, **Zhihan Xiong***, Maryam Fazel, Lalit Jain, Kevin Jamieson

- *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. [[paper](#)]

Parameterized Indexed Value Function for Efficient Exploration in Reinforcement Learning

- Tian Tan*, **Zhihan Xiong***, Vikranth R. Dwaracherla

- *Association for the Advancement of Artificial Intelligence (AAAI, Oral)*, 2020. [[arXiv](#)]

Work/ Research Experience

Meta

Visiting Researcher, FAIR Labs

supervised by Dr. **Lin Xiao**

Seattle, WA

Oct 2022–present

- Conducted research in theory of reinforcement learning and policy gradient algorithms.
- Proved the convergence of Soft Actor-Critic (SAC) algorithm under general function approximation in MDPs with continuous state-action space.

Bytedance

Research Scientist Intern, Applied Machine Learning (AML) Group

supervised by Dr. **Yingxiang Yang** and Dr. **Chong Wang**

Seattle, WA

Jun 2021–Sep 2021

- Explored utilizing data distribution’s periodicity through Fourier expansion for large-scale recommender systems.
- Published “*Fourier Learning with Cyclical Data*” at ICML 2022.

Zillow

Applied Scientist Intern, Personalization Team

supervised by Dr. **Luca Cazzanti**

Seattle, WA

Jun 2019–Sep 2019

- Participated in designing a new model with novel combination of decision tree and Thompson sampling to solve the recommendation-related contextual bandit problem.
- Developed an efficient method to accurately evaluate a policy using off-line data.

Teaching Experience

Graduate Teaching Assistant

CS 229: Machine Learning

CS 234: Reinforcement Learning

CS 229: Machine Learning

Stanford, CA

Autumn 2019

Winter 2020

Spring 2020

Awards & Honors

Meta AI Mentorship Program

NeurIPS 2022 Top Reviewers

ICML 2022 Travel Award

IFDS Research Assistantship

Yee Seung Ng Award

2022-2024

Autumn 2022

Summer 2022

Winter 2021

Spring 2017

Skills

Programming: Python (NumPy, Pandas, PyTorch), \LaTeX , R, C++

Languages: Chinese (native), English (fluent)