Runzhe Wu

2 W Loop Rd – New York – NY

rw646@cornell.edu • https://ziqian2000.github.io/

EDUCATION

Cornell Tech, Cornell University **New York, NY, USA**

Ph.D. Candidate, Computer Science *Aug. 2022 - Present*

- Advisor: Wen Sun. Committee Members: Wen Sun, Thorsten Joachims, Karthik Sridharan
- Started at Ithaca campus and then transferred to Cornell Tech campus at the start of my fourth year
- Anticipated graduation: Summer or Fall 2026

Shanghai Jiao Tong University **Shanghai, China**

B.E., Computer Science *Sep. 2018 - Jun. 2022*

- Member of ACM Honors Class (An Elite CS Program)
- Graduated with Highest Honor

RESEARCH INTERESTS

My research centers on **reinforcement learning**, focusing on both its core theories and its interplay with language models.

EXPERIENCE

Meta **New York, NY, USA**

Research Scientist Intern *May. 2025 - Aug. 2025*

- Manager: Yonathan Efroni
- Research on RL post-training of language models, with an emphasis on multi-task learning.

Northwestern University **Evanston, IL, USA**

Research Intern *Feb. 2021 - Feb. 2022*

- Advised by Zhaoran Wang and Zhuoran Yang.
- Research on reinforcement learning theory, particularly offline RL.

Shanghai Jiao Tong University **Shanghai, China**

Undergraduate Researcher *Aug. 2020 - Jun. 2021*

- Advised by Weinan Zhang and Yong Yu.
- Research on multi-agent reinforcement learning and population-based methods.

PUBLICATIONS & PREPRINTS

Imbalanced Gradients in RL Post-Training of Multi-Task LLMs

Runzhe Wu, Ankur Samanta, Ayush Jain, Scott Fujimoto, Jeongyeol Kwon, Ben Kretzu, Youliang Yu, Kaveh Hassani, Boris Vidolov, Yonathan Efroni

Preprint, 2025

Internalizing Self-Consistency in Language Models: Multi-Agent Consensus Alignment

Ankur Samanta, Akshayaa Magesh, Youliang Yu, **Runzhe Wu**, Ayush Jain, Daniel Jiang, Boris Vidolov, Paul Sajda, Yonathan Efroni, Kaveh Hassani

Preprint, 2025

Computationally Efficient RL under Linear Bellman Completeness for Deterministic Dynamics

Runzhe Wu*, Ayush Sekhari*, Akshay Krishnamurthy, Wen Sun

ICLR, 2025 **(Oral Presentation – Top 1.8%)**

Diffusing States and Matching Scores: A New Framework for Imitation Learning
Runzhe Wu, Yiding Chen, Gokul Swamy, Kianté Brantley, Wen Sun
ICLR, 2025

Making RL with Preference-based Feedback Efficient via Randomization
Runzhe Wu, Wen Sun
ICLR, 2024

Contextual Bandits and Imitation Learning via Preference-Based Active Queries
(Alphabetical order) Ayush Sekhari, Karthik Sridharan, Wen Sun, **Runzhe Wu**
NeurIPS, 2023

Selective Sampling and Imitation Learning via Online Regression
(Alphabetical order) Ayush Sekhari, Karthik Sridharan, Wen Sun, **Runzhe Wu**
NeurIPS, 2023

The Benefits of Being Distributional: Small-Loss Bounds for Reinforcement Learning
Kaiwen Wang, Kevin Zhou, **Runzhe Wu**, Nathan Kallus, Wen Sun
NeurIPS, 2023

Distributional Offline Policy Evaluation with Predictive Error Guarantees
Runzhe Wu, Masatoshi Uehara, Wen Sun
ICML, 2023

MALib: A Parallel Framework for Population-based Multi-agent Reinforcement Learning
Ming Zhou, Ziyu Wan, Hanjing Wang, Muning Wen, **Runzhe Wu**, Ying Wen, Yaodong Yang, Weinan Zhang, Jun Wang
JMLR, 2023

Offline Constrained Multi-Objective Reinforcement Learning via Pessimistic Dual Value Iteration
Runzhe Wu, Yufeng Zhang, Zhuoran Yang, Zhaoran Wang
NeurIPS, 2021

INVITED TALKS

Computationally Efficient RL under Linear Bellman Completeness for Deterministic Dynamics
ICLR 2025, Oral Session *Apr. 2025*

Computationally Efficient RL under Linear Bellman Completeness for Deterministic Dynamics
RL Theory Seminars *Nov. 2024*

AWARDS

Programming Contest Awards

- Gold Medal in the 2019 ICPC China Nanchang National Invitational Programming Contest
- Gold Medal in the 2018 ICPC Asia Hanoi Regional Programming Contest
- Gold Medal in the 2018 ICPC Asia Xuzhou Regional Programming Contest
- Gold Medal in the 2018 CCPC Qinhuangdao Regional Programming Contest
- Silver Medal in the 34th China National Olympiad in Informatics

Scholarships & Honors

- Outstanding Graduate of Shanghai *2022*
- Huawei Scholarship *2020*
- Rong Chang Scholarship *2019 & 2021*
- Zhiyuan Honorary Scholarship *2018 - 2020*

TEACHING EXPERIENCE

CS 4789: Introduction to Reinforcement Learning — **Cornell University**
Teaching Assistant — *Jan. 2023 - May. 2023*

CS 4700: Foundations of Artificial Intelligence — **Cornell University**
Teaching Assistant — *Aug. 2022 - Dec. 2022*

CS 147: Data Structure — **Shanghai Jiao Tong University**
Teaching Assistant — *Feb. 2022 - Jun. 2022*

CS 151: Programming — **Shanghai Jiao Tong University**
Teaching Assistant — *Sep. 2019 - Jan. 2020 & Sep. 2021 - Jan. 2022*

SERVICES

- **Reviewer** of ICML 2023-2025, NeurIPS 2023-2025, ICLR 2024-2026, COLT 2024
- **Mentor** of Cornell CS PhD Mentor Program — *2023 - 2024*
- **Organizer** of the 2021 ACM-Class Academic Symposium — *Jun. 2021*
- **Organizer** of the 2021 ACM-Class Sports Festival — *Apr. 2021 - May. 2021*
- **President** of ACM Honors Class, Shanghai Jiao Tong University — *Sep. 2020 - Jun. 2022*

SKILLS

Language English, Chinese
Programming Python, C/C++, LaTeX