

Yuanhao Wang | Curriculum Vitae

✉ yuanhao@princeton.edu • 🌐 www.cs.princeton.edu/~yuanhao

Education

- **Princeton University** 2020 –
 - Ph.D. Student in Computer Science, GPA 3.95/4.0
- **Tsinghua University** 2016 – 2020
 - B.Eng. in Computer Science, GPA 3.92/4.0
 - Enrolled in Special CS Pilot class supervised by Andrew Yao

Selected Publications

Yuanhao Wang*, Qinghua Liu*, and Chi Jin. Is RLHF More Difficult than Standard RL? *Advances in Neural Information Processing Systems*, 36, 2023.

Qinghua Liu*, **Yuanhao Wang***, and Chi Jin. Learning markov games with adversarial opponents: Efficient algorithms and fundamental limits. In *International Conference on Machine Learning*, pages 14036–14053. PMLR, 2022.

Chi Jin, Qinghua Liu, **Yuanhao Wang**, and Tiancheng Yu. V-Learning—A Simple, Efficient, Decentralized Algorithm for Multiagent RL. *Mathematics of Operations Research*, 2023.

Yuanhao Wang, Ruosong Wang, and Sham Kakade. An exponential lower bound for linearly realizable mdp with constant suboptimality gap. *Advances in Neural Information Processing Systems*, 34:9521–9533, 2021.

Yi Tian*, **Yuanhao Wang**, Tiancheng Yu*, and Suvrit Sra. Online learning in unknown markov games. In *International conference on machine learning*, pages 10279–10288. PMLR, 2021.

Guodong Zhang, **Yuanhao Wang**, Laurent Lessard, and Roger B Grosse. Near-optimal local convergence of alternating gradient descent-ascent for minimax optimization. In *International Conference on Artificial Intelligence and Statistics*, pages 7659–7679. PMLR, 2022.

Guodong Zhang and **Yuanhao Wang**. On the suboptimality of negative momentum for minimax optimization. In *International Conference on Artificial Intelligence and Statistics*, pages 2098–2106. PMLR, 2021.

Yuanhao Wang and Jian Li. Improved Algorithms for Convex-Concave Minimax Optimization. *Advances in Neural Information Processing Systems*, 33, 2020.

Yuanhao Wang*, Guodong Zhang*, and Jimmy Ba. On Solving Minimax Optimization Locally: A Follow-the-Ridge Approach. In *ICLR 2020*. OpenReview.net, 2020.

Yuanhao Wang*, Jiachen Hu*, Xiaoyu Chen, and Liwei Wang. Distributed Bandit Learning: Near-Optimal Regret with Efficient Communication. *arXiv preprint arXiv:1904.06309*, 2019.

Kefan Dong*, **Yuanhao Wang***, Xiaoyu Chen, and Liwei Wang. Q-learning with UCB Exploration is Sample Efficient for Infinite-Horizon MDP. In *ICLR 2020*. OpenReview.net, 2020.

Yuanhao Wang, Ying Li, Zhang-qi Yin, and Bei Zeng. 16-qubit IBM universal quantum computer can be fully entangled. *npj Quantum Information*, 4(1):46, 2018.

Professional Service

Reviewer for NeurIPS, ICML, ICLR, COLT, AISTATS and *Artificial Intelligence*.

Invited Talks

- **An exponential lower bound for linearly realizable MDP with constant suboptimality gap**, RL Theory Virtual Seminar, 2021
- **Is RLHF more difficult than standard RL?** University of Chicago, 2023

Selected Honors & Awards

- **1st Place Graduate Team**, Princeton COSCON, 2023
- **Francis Robbins Upton Fellowship in Engineering**, *awarded to outstanding incoming doctoral students in Princeton*, 2020
- **Outstanding Graduate Award**, Tsinghua University, 2020
- **National Scholarship in China**, *awarded to 1.5% students in Tsinghua*, 2018
- **9th of December Scholarship**, *among top 40 Tsinghua sophomores*, 2017
- **Gold Medal in Chinese Physics Olympiad**, *National Top 50*, 2015
- **Silver Medal in International Olympiad on Astronomy and Astrophysics**, *Global Top 15*, 2015

Skills

- **Coding:** Python, MATLAB, C++ & \LaTeX
- **Languages:** English (fluent); Chinese (native)

Other Experiences

- I have been TA for:
 - COS 445, 2023 Spring, Princeton University
 - ECE 434/COS 434, 2021 Fall, Princeton University
 - Mathematics for AI, 2020 Spring, Tsinghua University
- I visited the University of Toronto in Jul-Sep 2019 where I worked with Prof. Jimmy Ba.