Tianhao Wang **Last update: January 2026**

Contact Information

Halıcıoğlu Data Science Institute
University of California, San Diego
3234 Matthews Lane
La Jolla, CA 92093

203-305-9456
tianhaowang@ucsd.edu
https://tiiao.github.io
Google Scholar

Employment

Assistant Professor, Halıcıoğlu Data Science Institute, University of California, San Diego 2025 – present

Research Assistant Professor, Toyota Technological Institute at Chicago 2024 – 2025

Education

Ph.D. in Statistics and Data Science, Yale University 2024

- Advisor: Prof. Zhou Fan

M.A. in Statistics (en route), Yale University 2021

B.S. in Mathematics, University of Science and Technology of China (USTC) 2014 – 2018

- B.E. in Computer Science (dual)

Publications and Preprints

(*: equal contribution)

1. Shuo Xie*, Tianhao Wang*, Beining Wu, and Zhiyuan Li. "A Tale of Two Geometries: Adaptive Optimizers and Non-Euclidean Descent". *International Conference on Learning Representations (ICLR)*, 2026. Available at arXiv: 2511.20584.

2. Siyu Chen, Heejune Sheen, Xuyuan Xiong, Tianhao Wang, and Zhuoran Yang. "Taming Polysemanticity in LLMs: Provable Feature Recovery via Sparse Autoencoders". *International Conference on Learning Representations (ICLR)*, 2026. Available at arXiv: 2506.14002.

3. Xinhe Wang, Jin Huang, Xingjian Zhang, Tianhao Wang, Jiaqi W. Ma. "Your Reasoning Benchmark May Not Test Reasoning: Revealing Perception Bottleneck in Abstract Reasoning Benchmarks". Available at arXiv: 2512.21329.

4. Mohamad Amin Mohamadi, Tianhao Wang, and Zhiyuan Li. "Honesty over Accuracy: Trustworthy Language Models through Reinforced Hesitation". Available at arXiv: 2511.11500.

5. Xingjian Zhang, Tianhong Gao, Suliang Jin, Tianhao Wang, Teng Ye, Eytan Adar, and Qiaozhu Mei. "Through the Judge's Eyes: Inferred Thinking Traces Improve Reliability of LLM Raters". Available at arXiv: 2510.25860.

6. Robin Yadav, Shuo Xie, Tianhao Wang, and Zhiyuan Li. "Provable Benefit of Sign Descent: A Minimal Model Under Heavy-Tail Class Imbalance". NeurIPS 2025 Workshop on Optimization for Machine Learning (OPT 2025). **(Oral)**.

7. Max Lovig, Tianhao Wang, and Zhou Fan. "On Universality of Non-Separable Approximate Message Passing Algorithms". Available at arXiv: 2506.23010.

8. Shuo Xie, Tianhao Wang, Sashank Reddi, Sanjiv Kumar, and Zhiyuan Li. "Structured Preconditioners in Adaptive Optimization: A Unified Analysis". *International Conference on Machine Learning (ICML)*, 2025. Available at arXiv: 2503.10537.

9. Siyu Chen*, Beining Wu*, Miao Lu, Zhuoran Yang, and Tianhao Wang. "Can Neural Networks Achieve Optimal Computational-statistical Tradeoff? An Analysis on Single-Index Model". *International Conference on Learning Representations (ICLR)*, 2025 **(Oral)**.

10. Angeliki Giannou, Liu Yang, Tianhao Wang, Dimitris Papailiopoulos, and Jason D. Lee. "How well can Transformers emulate in-context Newton's method?" *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2025. Available at arXiv:2403.03183.

11. Siyu Chen*, Heejune Sheen*, Tianhao Wang, and Zhuoran Yang. "Unveiling Induction Heads: Provable Training Dynamics and Feature Learning in Transformers". *Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2024. Available at arXiv:2409:10559.

12. Xinyi Zhong*, Tianhao Wang*, and Zhou Fan. "Approximate Message Passing for orthogonally invariant ensembles: Multivariate non-linearities and spectral initialization". *Information and Inference*, 2024. Available at arXiv:2110.02318.

13. Tianhao Wang, Xinyi Zhong, and Zhou Fan. "Universality of Approximate Message Passing algorithms and tensor networks". *Annals of Applied Probability*, 2024. Available at arXiv:2206.13037.

14. Heejune Sheen, Siyu Chen, Tianhao Wang, and Harrison H. Zhou. "Implicit Regularization of Gradient Flow on One-Layer Softmax Attention". Available at arXiv:2403.08699.

15. Siyu Chen, Heejune Sheen, Tianhao Wang, and Zhuoran Yang. "Training dynamics of multi-head softmax attention for in-context learning: emergence, convergence, and optimality". *Conference on Learning Theory (COLT)*, 2024. Available at arXiv:2402.19442.

16. Zhou Fan, Roy R. Lederman, Yi Sun, Tianhao Wang, Sheng Xu. ''Maximum likelihood for high-noise group orbit estimation and single-particle cryo-EM". *Annals of Statistics*, 2024. Available at arXiv:2107.01305.

17. Runzhe Wang, Sadhika Malladi, Tianhao Wang, Kaifeng Lyu, and Zhiyuan Li. "The Marginal Value of Momentum for Small Learning Rate SGD". *International Conference on Learning Representations (ICLR)*, 2024. Available at arXiv:2307.15196.

18. Zhou Fan, Yi Sun, Tianhao Wang, Yihong Wu. "Likelihood landscape and maximum likelihood estimation for the discrete orbit recovery model". *Communications on Pure and Applied Mathematics*, 2023. Available at arXiv:2004.00041.

19. Ruitu Xu, Yifei Min, and Tianhao Wang. "Noise-adaptive Thompson sampling for linear contextual bandits". *Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2023.

20. Yifei Min, Jiafan He, Tianhao Wang, Quanquan Gu. "Cooperative multi-agent reinforcement learning: Asynchronous communication and linear function approximation". *International Conference on Machine Learning (ICML), 2023*. Available at arXiv:2305.06446.

21. Ruitu Xu, Yifei Min, Tianhao Wang, Michael I. Jordan, Zhaoran Wang, Zhuoran Yang. "Finding regularized competitive equilibria of heterogeneous agent macroeconomic models via reinforcement learning". *International Conference on Artificial Intelligence and Statistics (AISTATS), 2023*. Available at arXiv:2303.04833.

22. Zhiyuan Li, Tianhao Wang, Dingli Yu. "Fast mixing of stochastic gradient descent with normalization and weight decay". *Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2022.

23. Zhiyuan Li*, Tianhao Wang*, Jason D. Lee, Sanjeev Arora. "Implicit bias of gradient descent on reparametrized models: on equivalence to mirror descent". *Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2022. Available at arXiv:2207.04036.

24. Jiafan He*, Tianhao Wang*, Yifei Min*, Quanquan Gu. "A simple and provably efficient algorithm for asynchronous federeted linear bandits". *Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2022. Available at arXiv:2207.03106.

25. Yifei Min, Tianhao Wang, Ruitu Xu, Zhaoran Wang, Michael I. Jordan, Zhuoran Yang. "Learn to match with no regret: Reinforcement learning in Markov matching market". *Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2022 **(Oral)**. Available at arXiv:2203.03684.

26. Yifei Min, Jiafan He, Tianhao Wang, Quanquan Gu. "Learning stochastic shortest path with linear function approximation". *International Conference on Machine Learning (ICML)*, 2022. Available at arXiv:2110.12727.

27. Zhiyuan Li, Tianhao Wang, Sanjeev Arora. "What happens after SGD reaches zero loss? – A mathematical framework". *International Conference on Learning Representations (ICLR)*, 2022 **(Spotlight)**. Available at arXiv:2110.06914.

28. Pamela L Valentino, Tianhao Wang, Veronika Shabanova, Vicky Lee Ng, John C Bucuvalas, Amy G Feldman, Regino P Gonzalez-Peralta, Nitika Arora Gupta, Tamir A Miloh, Saeed Mohammad, Erika Pace, Shikha S Sundaram, Nada A Yazigi, Kyle Soltys, Society of Pediatric Liver Transplantation (SPLIT). "North American biliary stricture management strategies in children post liver transplant: multicenter analysis from the SPLIT Registry". *Liver Transplatation*, 2021.

29. Yifei Min*, Tianhao Wang*, Dongruo Zhou, Quanquan Gu. "Variance-aware off-policy evaluation with linear function approximation". *Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2021. Available at arXiv:2106.11960.

30. Tianhao Wang*, Dongruo Zhou*, Quanquan Gu. "Provably efficient reinforcement learning with linear function approximation under adaptivity constraints". *Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2021. Available at arXiv:2101.02195.

31. Pan Xu*, Tianhao Wang*, Quanquan Gu. "Continuous and discrete-time accelerated stochstic mirror descent for strongly convex functions". *International Conference on Machine Learning (ICML)*, 2018.

32. Pan Xu*, Tianhao Wang*, Quanquan Gu. "Accelerated stochastic mirror descent: From continuous-time dynamics to discrete-time algorithms". *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2018.

Honors and Awards

- Leonard F. Savage Prize for best written work — Yale, 2024
- Student Poster Competition Award — Rutgers University, 2023
  - At Conference on Recent Advances in Statistics and Data Science
- NeurIPS 2022 top reviewer — NeurIPS, 2022
- Wedworth W. Clarke Fellowship — Yale, 2021

Invited Talks and Seminars

- International Conference on Continuous Optimization, Los Angeles — 07/2025
- Midwest Machine Learning Symposium, Chicago — 06/2025
- International Indian Statistical Association Conference, Lincoln — 06/2025
- Midwest Optimization & Statistical Learning Conference, Chicago — 05/2025
- Flatiron Institute, CCM ML seminar — 04/2025
- Research at TTIC — 11/2024
- FAI seminar (online) — 08/2024
- Joint Statistical Meetings, Portland — 08/2024
- University of Notre Dame, ACMS Colloquium — 02/2024
- UIUC, Statistics Seminar — 02/2024
- Columbia University, Statistics Seminar — 01/2024
- UCSD, Halıcıoğlu Data Science Institute, Special Seminar Series — 01/2024
- UC Davis, Statistics Seminar — 01/2024
- UCSD, Department of Mathematics Colloquium — 12/2023
- INFORMS Annual Meeting, Phoenix — 10/2023

Teaching

Instructor, University of California, San Diego

- DSC 242: High-dimensional Probability and Statistics — Fall 2025

Teaching assistant, Yale University

- High-Dimensional Phenomena in Statistics and Learning — Spring 2023
- Intermediate Machine Learning — Spring 2022
- Statistical Inference — Fall 2020, Fall 2021
- Information Theory — Spring 2021
- Probability and Statistics — Fall 2019
- Stochastic Processes — Spring 2019, Spring 2020

Professional Service

Organizing committee for

- Midwest Optimization & Statistical Learning Conference 2025

Conference area chair for

- NeurIPS 2025
- ICLR 2025, 2026

Journal reviewer for

- Annals of Statistics
- IEEE Transactions on Information Theory
- Journal of the American Statistical Association
- Journal of Machine Learning Research
- Journal of the Royal Statistical Society: Series B
- Random Matrices: Theory and Applications
- SIAM Journal on Mathematics of Data Science
- Statistical Science

Conference reviewer for

- AAAI 2025
- AISTATS 2022, 2023, 2024
- COLT 2025
- ICLR 2023, 2024, 2025
- ICML 2022, 2023, 2024, 2025, 2026
- IEEE ITW 2023
- NeurIPS 2022, 2023, 2024
- Reinforcement Learning Conference 2024, 2025