Tianhao Wang

CONTACT INFORMATION

Toyota Technological Institute at Chicago
6045 South Kenwood Ave
Chicago, IL 60637
Phone: 203-305-9456

✉ tianhao.wang@ttic.edu
🖥 https://tianhaowang.ttic.edu
🔗 Linkedin
🎓 Google Scholar

EMPLOYMENT   Research Assistant Professor, Toyota Technological Institute at Chicago   2024 – present

EDUCATION   Ph.D. in Statistics and Data Science, Yale University   2018 – 2024

- Advisor: Prof. Zhou Fan

M.A. in Statistics (en route), Yale University   2018 – 2021

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

- B.E. in Computer Science (dual)

EXPERIENCE   Student researcher, Google DeepMind   09/2023 – 11/2023

Research intern, Google DeepMind   06/2023 – 08/2023

PUBLICATIONS AND PREPRINTS   (*: equal contribution)

1. Siyu Chen, Heejune Sheen, Tianhao Wang, and Zhuoran Yang. "Unveiling Induction Heads: Provable Training Dynamics and Feature Learning in Transformers". Available at arXiv:2409.10559.

2. 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.

3. Angeliki Giannou, Liu Yang, Tianhao Wang, Dimitris Papailiopoulos, and Jason D. Lee. "How well can Transformers emulate in-context Newton's method?" Available at arXiv:2403.03183.

4. 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.

5. 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.

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

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

8. 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.

9. 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.

10. Ruitu Xu, Yifei Min, and Tianhao Wang. "Noise-adaptive Thompson sampling for linear contextual bandits". *In Advances in Neural Information Processing Systems (NeurIPS)*, 2023.

11. 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.

12. 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.

13. Zhiyuan Li, Tianhao Wang, Dingli Yu. "Fast mixing of stochastic gradient descent with normalization and weight decay". *In Advances in Neural Information Processing Systems (NeurIPS)*, 2022.

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

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

16. 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". *In Advances in Neural Information Processing Systems (NeurIPS)*, 2022 **(Oral)**. Available at arXiv:2203.03684.

17. 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.

18. 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.

19. 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.

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

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

22. 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.

23. 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

- Conference Travel Fellowship   Yale, 2023

- 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

- ICML 2018 travel award   ICML, 2018

- Huang Yu Memorial Scholarship   USTC, 2017

INVITED TALKS AND SEMINARS		
	• FAI seminar	08/2024
	• Joint Statistical Meetings	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	10/2023

- – Universality of Approximate Message Passing algorithms and tensor networks
  - International Conference on Machine Learning — 07/2022
    - – Learning stochastic shortest path with linear function approximation
    - – Implicit bias of gradient descent on reparameterized models

TEACHING EXPERIENCE

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

SERVICE

Journal reviewer for

- IEEE Transactions on Information Theory
- Journal of the American Statistical Association
- Random Matrices: Theory and Applications
- The Annals of Statistics

Conference reviewer for

- AISTATS 2022, 2023, 2024
- ICLR 2023, 2024
- ICML 2022, 2023, 2024
- IEEE ITW 2023
- NeurIPS 2022, 2023, 2024