

Song Mei

367 Evans Hall
Berkeley, CA, 94720

Website: <https://stat.berkeley.edu/~songmei/>
Email: songmei@berkeley.edu

Appointment

Assistant Professor, University of California, Berkeley

Department of Statistics. 07/2020 - present
Department of Electrical Engineering and Computer Sciences. 07/2021 - present

Education

Ph.D. Computational and Mathematical Engineering, Stanford University. 09/2014 - 07/2020
B.S. Mathematics, Peking University. 09/2010 - 06/2014

Research Interests

Foundations of deep learning and generative AI. Reinforcement learning. High dimensional statistics.

Grants and Awards

Amazon Research Award (Nominated). 2024.
Google Research Scholar Award. 2024.
NSF CAREER Award. 2024-2029.
NSF CCF-2315725. 2023-2026.
NSF DMS-2210827. 2022-2025.
Stanford Graduate Fellowship. 2014-2017.

Students supervised

Licong Lin: Stat PhD at UC Berkeley, 3rd year, co-advised with Peter Bartlett.
Nikhil Ghosh: Stat PhD at UC Berkeley, 5rd year, co-advised with Bin Yu.
Taejoo Ahn: Stat PhD at UC Berkeley, 6th year.
Zitong Yang: EECS master student at UC Berkeley, summer intern.
Ziang Song: Math undergrad student at Peking University, summer intern.
Fan Chen: Math undergrad student at Peking University, summer intern.

Teaching

University of California, Berkeley

Instructor, STAT210B. Theoretical Statistics. Spring 2023, Spring 2022.
Instructor, STAT210A. Theoretical Statistics. Fall 2022.
Instructor, STAT154. Modern Statistical Prediction and Machine Learning. Fall 2021. Spring 2024.
Instructor, STAT260. Mean Field Asymptotics in Statistical Learning. Spring 2021.

Professional Services

Seminar and workshop organization

July. 2020-Jan 2022. *OW Seminar, Mathematics of Machine Learning, seminar and workshop series.*

Sept. 2020-Oct 2022. *Neyman seminar, University of California, Berkeley.*

Editorial activity

Area Chair: COLT, 2022, 2023, 2024. COLM 2024.

Associate Editor: Journal of Machine Learning.

Publications

Citations ([Google Scholar](#)): 3459 in total. Date: Apr 7, 2024.

Journal Publications

1. Michael Celentano, Zhou Fan, and **Song Mei**. (2023) Local convexity of the TAP free energy and AMP convergence for Z2-synchronization. *The Annals of Statistics* 51 (2), 519-546.
2. **Song Mei**, Theodor Misiakiewicz, and Andrea Montanari. (2021) Generalization error of random features and kernel methods: hypercontractivity and kernel matrix concentration. *Applied and Computational Harmonic Analysis*, 59, 3-84.
3. Behrooz Ghorbani, **Song Mei**, Theodor Misiakiewicz, and Andrea Montanari. (2021) When do neural networks outperform kernel methods? *Journal of Statistical Mechanics: Theory and Experiment*.
4. **Song Mei**, and Andrea Montanari. (2021) The generalization error of random features regression: Precise asymptotics and double descent curve. *Communications on Pure and Applied Mathematics*, 75 (4), 667-766.
5. Zhou Fan, **Song Mei**, and Andrea Montanari. (2021) TAP free energy, spin glasses, and variational inference. *The Annals of Probability* 49 (1), 1-45.
6. Behrooz Ghorbani, **Song Mei**, Theodor Misiakiewicz, and Andrea Montanari. (2020) Linearized two-layers neural networks in high dimension. *The Annals of Statistics* 2021, Vol. 49, No. 2, 1029-1054.
7. Behrooz Ghorbani, **Song Mei**, Theodor Misiakiewicz, and Andrea Montanari. (2019) Discussion of “Nonparametric Regression using Deep Neural Networks with ReLU Activation Function”. *The Annals of Statistics* 48 (4), 1898-1901.
8. Gerard Ben Arous, **Song Mei**, Andrea Montanari, and Mihai Nica. (2019) The landscape of the spiked tensor model. *Communications on Pure and Applied Mathematics* 72 (11), 2282-2330.
9. **Song Mei**, Andrea Montanari, and Phan-Minh Nguyen. (2018) A mean field view of the landscape of two-layers neural network. *Proceedings of the National Academy of Sciences* 115, E7665-E7671.
10. **Song Mei**, Yu Bai, and Andrea Montanari. (2018) The landscape of empirical risk for non-convex losses. *The Annals of Statistics* 46 (6A), 2747-2774.
11. **Song Mei**, and Pingwen Zhang. (2015) On a molecular based Q-tensor model for liquid crystals with density variations. *Multiscale Modeling and Simulation* 13 (3), 977-1000.

Conference Publications

1. Fan Chen, Huan Wang, Caiming Xiong, **Song Mei**, and Yu Bai. (2023) Lower Bounds for Learning in Revealing POMDPs. *International Conference on Machine Learning (ICML)*, 2023.
2. Fan Chen, Yu Bai, and **Song Mei**. (2022) Partially Observable RL with B-Stability: Unified Structural Condition and Sharp Sample-Efficient Algorithms. *The International Conference on Learning Representations (ICLR)*, 2023.
3. Yu Bai, Chi Jin, **Song Mei**, Ziang Song, and Tiancheng Yu. (2022) Efficient Phi-Regret Minimization in Extensive-Form Games via Online Mirror Descent. *Neural Information Processing Systems (NeurIPS)*, 2022.
4. Ziang Song, **Song Mei**, and Yu Bai. (2022) Sample-Efficient Learning of Correlated Equilibria in Extensive-Form Games. *Neural Information Processing Systems (NeurIPS)*, 2022.
5. Theodor Misiakiewicz, and **Song Mei**. (2021) Learning with convolution and pooling operations in kernel methods. *Neural Information Processing Systems (NeurIPS)*, 2022.
6. Yu Bai, Chi Jin, **Song Mei**, and Tiancheng Yu. (2022) Near-Optimal Learning of Extensive-Form Games with Imperfect Information *International Conference on Machine Learning (ICML)*, 2022.
7. Yu Bai, **Song Mei**, Huan Wang, Yingbo Zhou, and Caiming Xiong. (2021) Efficient and Differentiable Conformal Prediction with General Function Classes. *The International Conference on Learning Representations (ICLR)*, 2022.
8. Nikhil Ghosh, **Song Mei**, Bin Yu. (2021) The Three Stages of Learning Dynamics in High-Dimensional Kernel Methods. *The International Conference on Learning Representations (ICLR)*, 2022.
9. Ziang Song, **Song Mei**, Yu Bai. (2021) When Can We Learn General-Sum Markov Games with a Large Number of Players Sample-Efficiently? *The International Conference on Learning Representations (ICLR)*, 2022.
10. Yu Bai, **Song Mei**, Huan Wang, and Caiming Xiong. (2021) Understanding the Under-Coverage Bias in Uncertainty Estimation. *Neural Information Processing Systems (NeurIPS)*, 2021.
11. Zitong Yang, Yu Bai, and **Song Mei**. (2021) Exact Gap between Generalization Error and Uniform Convergence in Random Feature Models. *International Conference on Machine Learning (ICML)*, 2021.
12. Yu Bai, **Song Mei**, Huan Wang, and Caiming Xiong. (2021) Don't Just Blame Over-parametrization for Over-confidence: Theoretical Analysis of Calibration in Binary Classification. *International Conference on Machine Learning (ICML)*, 2021.
13. **Song Mei**, Theodor Misiakiewicz, and Andrea Montanari. (2021) Learning with invariances in random features and kernel models. *Conference of Learning Theory (COLT)*, 2021.
14. Behrooz Ghorbani, **Song Mei**, Theodor Misiakiewicz, and Andrea Montanari. (2020) When do neural networks outperform kernel methods? *Neural Information Processing Systems (NeurIPS)*, 2020.
15. Behrooz Ghorbani, **Song Mei**, Theodor Misiakiewicz, and Andrea Montanari. (2019) Limitations of lazy training of two-layers neural networks. *Neural Information Processing Systems (NeurIPS)*, 2019.

16. **Song Mei**, Theodor Misiakiewicz, and Andrea Montanari. (2019) Mean-field theory of two-layers neural networks: dimension-free bounds and kernel limit. *Conference of Learning Theory (COLT)*, 2019.
17. **Song Mei**, Theodor Misiakiewicz, Andrea Montanari, and Roberto I. Oliveira. (2017) Solving SDPs for synchronization and MaxCut problems via the Grothendieck inequality. *Conference of Learning Theory (COLT)*, 2017.
18. Yu Bai, Fan Chen, Huan Wang, Caiming Xiong, and **Song Mei**. (2023) Transformers as Statisticians: Provable In-Context Learning with In-Context Algorithm Selection. *NeurIPS 2023 (Spotlight)*.
19. Hengyu Fu, Tianyu Guo, Yu Bai, and **Song Mei**. (2023) What can a Single Attention Layer Learn? A Study Through the Random Features Lens. *NeurIPS 2023*.
20. Tianyu Guo, Wei Hu, **Song Mei**, Huan Wang, Caiming Xiong, Silvio Savarese, Yu Bai. (2023) How Do Transformers Learn In-Context Beyond Simple Functions? A Case Study on Learning with Representations. *International Conference on Learning Representations (ICLR)*, 2024.
21. Licong Lin, Yu Bai, and **Song Mei**. (2023) Transformers as Decision Makers: Provable In-Context Reinforcement Learning via Supervised Pretraining. *International Conference on Learning Representations (ICLR)*, 2024.

Preprints and Submissions

1. Ruiqi Zhang, Licong Lin, Yu Bai, and **Song Mei**. (2024) Negative Preference Optimization: From Catastrophic Collapse to Effective Unlearning. *arXiv preprint arXiv:2404.05868*.
2. Taejoo Ahn, Licong Lin, and **Song Mei**. (2022) Near-optimal multiple testing in Bayesian linear models with finite-sample FDR control. *arXiv preprint arXiv:2211.02778*.
3. Fan Chen, **Song Mei**, and Yu Bai. (2022) Unified Algorithms for RL with Decision-Estimation Coefficients: No-Regret, PAC, and Reward-Free Learning. *arXiv preprint arXiv:2209.11745*.
4. **Song Mei**, and Yuchen Wu. (2023) Deep Networks as Denoising Algorithms: Sample-Efficient Learning of Diffusion Models in High-Dimensional Graphical Models. *arXiv preprint arXiv:2309.11420*.
5. Hui Xu, **Song Mei**, Stephen Bates, Jonathan Taylor, and Robert Tibshirani (2023) Uncertainty Intervals for Prediction Errors in Time Series Forecasting. *arXiv preprint arXiv:2309.07435*.
6. Michael Celentano, Zhou Fan, Licong Lin, and **Song Mei** (2023) Mean-field variational inference with the TAP free energy: Geometric and statistical properties in linear models. *arXiv preprint arXiv:2311.08442*.
7. Leo Zhou, Joao Basso, and **Song Mei** (2024) Statistical Estimation in the Spiked Tensor Model via the Quantum Approximate Optimization Algorithm. *arXiv preprint arXiv:2402.19456*.