Eric Zhao

Homepage: eric-zhao.com
Email: eric28zh@gmail.com

EXPERIENCE

OpenAI*: Member of Technical Staff* — July 2025 - Present
I'm a research scientist working on LLMs and reasoning.

EDUCATION

UC Berkeley — 2021-2025
Ph.D., Computer Science
Advisors: Michael I. Jordan & Nika Haghtalab.

Caltech (California Institute of Technology) — 2016-2020
B.S., Computer Science
Advisors: Adam Wierman & Yisong Yue.

AWARDS

Google PhD Fellowship, Machine Intelligence (2024-2025)
NSF Graduate Research Fellowship (2023-2026)
NeurIPS Best Paper Award (2022)

INTERNSHIPS

Google Research*: Research Intern* — May 2023 - May 2025
Salesforce Research*: Research Intern* — June - Sept. 2020
Uber ATG*: Software Engineering Intern* — June - Sept. 2019
Google*: Software Engineering Intern* — April - June 2019

PUBLICATIONS

(α-β) denotes when authors are ordered alphabetically.

[Link] Eric Zhao, Pranjal Awasthi, Nika Haghtalab. "*From Style to Facts: The Boundaries of Knowledge Injection with Finetuning*," Feb. 2025.. In the 39th Annual Conference on Neural Information Processing Systems (NeurIPS).

[Link] (α-β) Mingda Qiao, Eric Zhao. "*Truthfulness of Decision-Theoretic Calibration Measures*," Jan 2025. In 38th Annual Conference on Learning Theory (COLT).

[Link] Eric Zhao, Pranjal Awasthi, Sreenivas Gollapudi. "*Sample, Scrutinize and Scale: Effective Inference-Time Search by Scaling Verification*," Jan. 2025. In the 42nd International Conference on Machine Learning (ICML).

[Link] (α-β) Jessica Dai, Nika Haghtalab, Eric Zhao. "*Learning With Multi-Group Guarantees For Clusterable Subpopulations*," Nov. 2024. In the 42nd International Conference on Machine Learning (ICML).

[Link] (α-β) Nika Haghtalab, Mingda Qiao, Kunhe Yang, Eric Zhao. "*Truthfulness of Calibration Measures*," May. 2024. In the 38th Annual Conference on Neural Information Processing Systems (NeurIPS).

[Link] Eric Zhao, Pranjal Awasthi, Zhengdao Chen, Sreenivas Gollapudi, Daniel Delling. "*Semantic Routing via Autoregressive Modeling*," May. 2024. In the 38th Annual Conference on Neural Information Processing Systems (NeurIPS).

[Link] (α-β) Pranjal Awasthi, Nika Haghtalab, Eric Zhao. "*Open Problem: The Sample Complexity of Multi-Distribution Learning for VC Classes*," May. 2023. In Proceedings of the 36th Annual Conference on Learning Theory (COLT).

[Link] (α-β) Nika Haghtalab, Michael I. Jordan, Eric Zhao. "*A Unifying Perspective on Multicalibration: Game Dynamics for Multi-Objective Learning*," Feb. 2023. In Proceedings of the 37th Annual Conference on Neural Information Processing Systems (NeurIPS).

[Link] (α-β) Nika Haghtalab, Michael I. Jordan, Eric Zhao. "*On-Demand Sampling: Learning Optimally from Multiple Distributions*," May. 2022. In Proceedings of the 36th Annual Conference on Neural Information Processing Systems (NeurIPS). NeurIPS Best Paper Award.

[Link] Eric Zhao, Alex Trott, Caiming Xiong, Stephan Zheng. "*Learning to Play General-Sum Games Against Multiple Boundedly Rational Agents*," Jan. 2021. In Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI).

[Link] Eric Zhao, Anqi Liu, Animashree Anandkumar, Yisong Yue. "*Active Learning under Label Shift*," April 2021. In Proceedings of The 24th International Conference on Artificial Intelligence and Statistics (AISTATS). Also appeared at the ICML 2020 Workshop on Real World Experiment Design and Active Learning.

PREPRINTS

[Link] Eric Zhao, Jessica Dai, Pranjal Awasthi. "*The Limits of Preference Data for Post-Training*," May. 2025.

[Link] (α-β) Naman Agarwal, Pranjal Awasthi, Satyen Kale, Eric Zhao. "*Stacking as Accelerated Gradient Descent*," Nov. 2024.

[Link] (α-β) Emilio Calvano, Nika Haghtalab, Ellen Vitercik, Eric Zhao. "*Algorithmic Content Selection and the Impact of User Disengagement*," Nov. 2024.

TALKS

ACM Conference on Economics and Computation (EC 2025). *"Tutorial on Decision-Theoretic Forecasting"*

European Research Council OCEAN Seminar (2025). *"Game Dynamics for Multi-Objective Learning: A Modular Approach to Calibration and Multicalibration."*

Symposium on Foundations of Responsible Computing (2024). *"A Unifying Perspective on Multicalibration: Game Dynamics for Multi-Objective Learning."*

Google GNN Summit (2024). *"Semantic Routing via Auto-regressive Modeling."*

INFORMS Annual Meeting, Revenue Management and Pricing Session (2024). *"Algorithmic Content Selection and the Impact of User Disengagement."*

Symposium on Foundations of Responsible Computing (2023). *"On-Demand Sampling: Learning Optimally from Multiple Distributions."*

Conference on Learning Theory, Open Problems (COLT 2023). *"The Sample Complexity of Multi-Distribution Learning for VC Classes."*

INFORMS Annual Meeting, Frontiers of Algorithmic Fairness Session (2023). *"Multi-Objective Learning: A Unifying Perspective On Fairness."*

Conference on Neural Information Processing Systems (NeurIPS 2022). *"On-Demand Sampling: Learning Optimally from Multiple Distributions."*

TEACHING

Berkeley grad TA for *CS 170: Efficient Algorithms and Intractable Problems* (Nika Haghtalab, John Wright 2025), *CS 227: Convex Optimization* (Ben Recht 2024).

Caltech undergrad TA for *CS 144: Algorithmic Economics* (Adam Wierman 2019) and *CS 155: Machine Learning* (Yisong Yue 2019).

SERVICE

Journal reviewer for Journal of Machine Learning Research (JMLR).

Conference reviewer for International Conference on Machine Learning (ICML, Outstanding Reviewer Award), International Conference on Learning Representations (ICLR), Neural Information Processing Systems (NeurIPS), International Conference on Artificial Intelligence and Statistics (AISTATS), Annual Conference on Learning Theory (COLT), ACM Conference on Economics & Computation.