

Eric Zhao

Homepage: www.eric-zhao.com

Email: eric.zh@berkeley.edu

SUMMARY

I'm a 4th year computer science PhD candidate at UC Berkeley, where I'm a member of the Berkeley Artificial Intelligence Research Lab (BAIR) and Berkeley Theory Group. My research develops algorithms and mathematical foundations for multi-objective machine learning. I am particularly interested in multi-objective learning as a means of (1) understanding the emergent capabilities of large language models and (2) designing algorithmic approaches for data collection, forecasting, and search.

EDUCATION

University of California, Berkeley 2021-2026

Ph.D., Computer Science. GPA 3.9/4.

Advised by Nika Haghtalab & Michael I. Jordan. Simons Institute visitor (2022, 2024).

California Institute of Technology 2016-2020

B.S., Computer Science. GPA 3.9/4.

Advised by Adam Wierman & Yisong Yue.

AWARDS

Google PhD Fellowship (2024)

NSF Graduate Research Fellowship (2023)

Neurips Outstanding Paper Award (2022)

PUBLICATIONS

[[Link](#)] Nika Haghtalab, Mingda Qiao, Kunhe Yang, **Eric Zhao** (Alphabetical author ordering). “*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](#)] Pranjal Awasthi, Nika Haghtalab, **Eric Zhao** (Alphabetical author ordering). “*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](#)] Nika Haghtalab, Michael I. Jordan, **Eric Zhao** (Alphabetical author ordering). “*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](#)] Nika Haghtalab, Michael I. Jordan, **Eric Zhao** (Alphabetical author ordering). “*On-Demand Sampling: Learning Optimally from Multiple Distributions*,” May. 2022. In Proceedings of the 36th Annual Conference on Neural Information Processing Systems (NeurIPS). *Outstanding 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.

WORKING PAPERS

[Link] Naman Agarwal, Pranjal Awasthi, Satyen Kale, **Eric Zhao** (Alphabetical author ordering). “*Stacking as Accelerated Gradient Descent*,” under review, Feb. 2024.

Emilio Calvano, Nika Haghtalab, Ellen Vitercik, **Eric Zhao** (Alphabetical author ordering). “*When Market Friction Increases User Engagement: Incentivizing Investment in Future Interactions*,” under review, Feb. 2024.

Eric Zhao, Tatjana Chavdarova, Michael I. Jordan. “*Learning Variational Inequalities: Generalization Bounds and Applications to Empirical Games*,” under review, Oct. 2023.

[Link] **Eric Zhao**, De-An Huang, Hao Liu, Zhiding Yu, Anqi Liu, Olga Russakovsky, Anima Anandkumar. “*Scaling Bias Mitigation with Multiple Fairness Tasks and Multiple Protected Attributes*,” technical report, Sept. 2021.

EXPERIENCE

Google Research: Student Researcher May. 2023 - Present

- Trained large language models (LLM) on customized routing tasks (Google Maps). Conducted theoretical analysis of the LLM training technique of “stacking”.

Salesforce Research: Research Intern June - Sept. 2020

- Ran deep reinforcement learning experiments on finding robust and near-optimal tax policies in simulated trade-and-barter economies.

Nvidia Research: Research Intern Feb. 2021 - June. 2021

- Ran computer vision experiments on flaws in machine learning bias mitigation methods and studied alternatives for intersectional & label-scarce settings.

Uber: Self-Driving Engineering Intern June - Sept. 2019

- Developed safety-critical computer vision algorithms for Uber’s self-driving cars, extending blockage localization range by up to 20% in live track tests.

Google: Software Engineering Intern April - June 2019

- Prototyped in-conversation error-handling capabilities for Dialogflow chatbots.

Caltech: Student Researcher Dec. 2018 - June 2020

- Assisted case study on California school choice lottery that boosted match rates 10%.

Bloomberg LP: Software Engineering Intern June - Aug. 2018

- Developed anomaly detection algorithms for financial data,

STARTUPS

HVF Labs: Visitor Nov. - Dec. 2018

- Helped pilot credit card reward program with Max Levchin in his venture incubator.

Whimply Inc.: Cofounder Oct. 2017 - Mar. 2018

- Built chatbots for eCommerce sites inc. Fortune 500s; in MassChallenge accelerator.

Brewgorithm: Cofounder June - Oct. 2017

- Cofounded startup building recommendation engines for beverages/alcohol. Joined a ZX-Ventures accelerator, raised seed round from AB InBev SA/NV.

SERVICE

Reviewer/Sub-Reviewer: AISTats 2021, ICML 2021, Neurips 2022, ICLR 2022, ICML 2022 (Outstanding Reviewer award), ICML 2023, ICML 2024.

Undergraduate Mentor: Mentored students in the Equal Access to Application Assistance (EAAA) Program and BAIR Undergraduate Mentorship Program. (2021-2023)

Volunteer Teacher: Served as a part-time computer science teacher at low income schools with Microsoft Philanthropy’s TEALS Program (2020-2021)