

Haofeng “Fred” Zhang

Email : z0@berkeley.edu

Website: <https://fredzhang.me/>

EDUCATION

- **University of California, Berkeley** Berkeley, CA
Ph.D. in Computer Science *September 2019 – May 2024*
 - Advisor: [Jelani Nelson](#)
- **Harvard University** Cambridge, MA
Ph.D. in Computer Science (transferred to UC Berkeley with advisor) *August 2018 – May 2019*
- **Duke University** Durham, NC
B.S. in Computer Science, with Highest Distinction *May 2018*
B.S. in Mathematics *May 2018*

SELECTED PUBLICATIONS

- [1] **[Approaching Human-Level Forecasting with Language Models](#)**
Danny Halawi*, [Fred Zhang*](#), Chen Yue-Han*, and Jacob Steinhardt (* Equal contribution)
Preprint.
- [2] **[Towards Best Practices of Activation Patching in Language Models: Metrics and Methods](#)**
[Fred Zhang](#) and Neel Nanda
International Conference on Learning Representations (ICLR), 2024.
- [3] **[Privately Estimating a Gaussian: Efficient, Robust and Optimal](#)**
Daniel Alabi*, Pravesh K. Kothari*, Prayaag Venkat*, Pranay Tankala* and Fred Zhang* (Alphabetical order)
ACM Symposium on Theory of Computing (STOC), 2023
- [4] **[Online Prediction in Sub-linear Space](#)**
Binghui Peng* and Fred Zhang* (Alphabetical order)
ACM-SIAM Symposium on Discrete Algorithms (SODA), 2023. * *Winner of Best Student Paper.*
- [5] **[Robust and Heavy-Tailed Mean Estimation Made Simple, via Regret Minimization](#)**
Samuel B. Hopkins*, Jerry Li* and Fred Zhang* (Alphabetical order)
Neural Information Processing Systems (NeurIPS), 2020.
- [6] **[A Fast Spectral Algorithm for Mean Estimation with Sub-Gaussian Rates](#)**
Zhixian Lei*, Kyle Luh*, Prayaag Venkat* and Fred Zhang* (Alphabetical order)
Conference on Learning Theory (COLT), 2020.

PROFESSIONAL EXPERIENCE

- **Google** New York City, NY
Research Intern *May 2023 – August 2023*
 - Research in clustering latent activations in neural networks and mechanistic interpretability of language model.
 - Mentored by [Matthew Fahrbach](#), [Neel Nanda](#) and [Peilin Zhong](#).
- **Google** Pittsburgh, PA
Research Intern *May 2022 – August 2022*
 - Research in sketching methods, efficient online learning and optimization at [Google Brain](#) (now DeepMind).
 - Mentored by [Richard Zhang](#) and [David P. Woodruff](#).
- **Google** Mountain View, CA
Software Engineering Intern *May 2016 – August 2016*
 - Applied machine learning and natural language processing at [Google Research](#).

ACADEMIC EXPERIENCE

- **Simons Institute for the Theory of Computing** Berkeley, CA
Visiting Graduate Student *August 2020 – December 2020*
 - Participated the Fall 2020 program on [Probability, Geometry, and Computation in High Dimensions](#).
- **Department of Computer Science, Duke University** Durham, NC
Undergraduate Research Fellow *May 2017 – August 2017*
 - Research in online algorithms for competitive caching, advised by [Rong Ge](#) and [Debmalya Panigrahi](#).

TEACHING EXPERIENCE

- **Department of Electrical Engineering and Computer Sciences, UC Berkeley** Berkeley, CA
Graduate Student Instructor *January 2020 – December 2020*
 - Teaching assistant for [CS 294-165: Sketching Algorithms](#) (Fall 2020).
 - Teaching assistant for [CS 170: Efficient Algorithms and Intractable Problems](#) (Spring 2020).
- **Department of Computer Science, Duke University** Durham, NC
Undergraduate Teaching Assistant *January 2016 – May 2018*
 - Teaching assistant for CompSci 330: Design and Analysis of Algorithms for 5 semesters: [Spring 2016](#), [Fall 2016](#), [Spring 2017](#), [Fall 2017](#) and [Spring 2018](#).

INVITED TALKS

- [1] Privately Estimating a Gaussian: Efficient, Robust and Optimal. *Algorithms Seminar, Google Research*, February, 2024.
- [2] Online Prediction in Sub-linear Space. *Algorithms Seminar, Google Research*, June, 2023.
- [3] Online Prediction in Sub-linear Space. *Algorithms Seminar, University of Illinois Urbana-Champaign*, April, 2023.
- [4] Online Prediction in Sub-linear Space. *Theory Seminar, Toyota Technological Institute at Chicago*, April, 2023.
- [5] Online Prediction in Sub-linear Space. *Theory Seminar, Northwestern University*, April, 2023.
- [6] Online Prediction in Sub-linear Space. *Theory Lunch, UC Berkeley*, March, 2023.
- [7] Optimal Robustness-Consistency Tradeoffs for Learning-Augmented Online Algorithms. *Foundations of Data Science Institute Retreat, Simons Institute for the Theory of Computing*, January, 2022.

SELECTED AWARDS

- Best Student Paper Award, ACM-SIAM Symposium on Discrete Algorithms (SODA 2023) 2023
- Outstanding Reviewer, Conference on Neural Information Processing Systems (NeurIPS 2022, 2023) 2022, 2023
- Karina A. Chen Graduate Student Research Fellowship, Harvard University 2018–2019
- Deans' Summer Research Fellowship, Duke University 2017

SKILLS

- **Programming Languages:** Python, C++, MATLAB, HTML/CSS
- **Machine Learning:** PyTorch, JAX, NumPy, TransformerLens (mechanistic interpretability)