

SPENCER FREI

sfrei@google.com ◇ spencerfrei.github.io ◇ Google Scholar

EXPERIENCE

Senior Research Scientist, Google DeepMind (Frontier AI)
Pretraining research in JAX on TPUs.

Jan. 2025–present

- Led a quantitative study of the sources of run-to-run variation in LLM pretraining, showing that floating point arithmetic order has nearly as large an effect as data and initialization seeds.
- Developed more reliable algorithms for pretraining scaling law estimation. Paper in preparation.
- Built agentic research workflow that facilitated multiple parallel multi-month research projects.
- Mentored a student researcher through a research project, from problem formulation to a paper now in preparation.
- Additional pretraining research not publicly disclosable.

Assistant Professor, Department of Statistics, University of California, Davis
Tenure-track faculty in Statistics, member of the Graduate Group in Computer Science. Long-term participant, Simons Institute Fall 2024 program on Modern Paradigms in Generalization.

Sep. 2023–Jan. 2025

Postdoctoral Fellow, Simons Institute for the Theory of Computing, UC Berkeley

2021–2023

Mentors: Peter Bartlett and Bin Yu. Work on benign overfitting, implicit regularization, and in-context learning in transformers. Visitor at EPFL (Sep.–Oct. 2022), hosted by Emmanuel Abbé.

Applied Scientist Intern, Amazon Alexa AI, Cambridge, MA
Natural language understanding with LLMs.

Summer 2020

SELECTED PUBLICATIONS AND WRITING

full list of 23 on Google Scholar

1. **S. Frei**. Quantifying pretraining variance: Data, initialization, and floating point arithmetic. *Technical blog post*, 2026. spencerfrei.github.io/blog/2026-08-17-pretraining-variance
2. Fangyuan Lin, **S. Frei**, Victor de la Peña. Bounded difference concentration for infinitely exchangeable sequences with applications to AI benchmark uncertainty. Preprint, arXiv:2606.17426.
3. **S. Frei**, Gal Vardi. Trained transformer classifiers generalize and exhibit benign overfitting in-context. *ICLR 2025*.
4. Usman Anwar, Johannes von Oswald, Louis Kirsch, David Krueger, **S. Frei**. Understanding in-context learning of linear models in transformers through an adversarial lens. *TMLR 2025*.
5. Ruiqi Zhang, **S. Frei**, Peter L. Bartlett. Trained transformers learn linear models in-context. *JMLR 2024*.
6. **S. Frei***, Gal Vardi*, Peter L. Bartlett, Nati Srebro. The double-edged sword of implicit bias: Generalization vs. robustness in ReLU networks. *NeurIPS 2023*.
7. **S. Frei***, Gal Vardi*, Peter L. Bartlett, Nati Srebro, Wei Hu. Implicit bias in leaky ReLU networks trained on high-dimensional data. *ICLR 2023 (Spotlight)*.
8. **S. Frei**, Niladri Chatterji, Peter L. Bartlett. Benign overfitting without linearity: Neural network classifiers trained by gradient descent for noisy linear data. *COLT 2022*.

* indicates equal contribution.

EDUCATION

Ph.D., Statistics , UCLA Co-advisors: Quanquan Gu and Ying Nian Wu.	<i>2015–2021</i>
M.Sc., Mathematics , University of British Columbia	<i>2013–2015</i>
B.Sc., Mathematics , McGill University First class honours.	<i>2009–2013</i>

HONORS

- Selected to present the NeurIPS 2023 Tutorial, “Reconsidering Overfitting in the Age of Overparameterized Models.”
- Rising Star in Machine Learning, University of Maryland, 2022.
- Best Reviewer Award: ICML 2020, ICML 2021, NeurIPS 2021, ICLR 2022.
- Long Talk, ICML 2021 (top 3.0% of 5,514 submissions).

INVITED TALKS

- 50+ invited talks and tutorials, at venues including Oxford, Cambridge, Stanford, Columbia, EPFL, ETH Zürich, Apple Machine Learning Research, Microsoft Research, and Google DeepMind.

MENTORSHIP

Student researchers, Google DeepMind

- Bruno Mlodozieniec (Cambridge)

PhD students, UC Berkeley and UC Davis

- Ruiqi Zhang (Berkeley; now at Citadel Securities)
- Nikhil Ghosh (Berkeley; now at the Flatiron Institute)
- Neil Mallinar (UC San Diego; now at Pryon)
- Austin Zane (Berkeley; now at Google)

PROFESSIONAL SERVICE

- Area Chair / Senior PC: NeurIPS 2023, ALT 2024, NeurIPS 2024, ALT 2025, ICML 2025, ICLR 2026; Theoretical Foundations of Foundation Models @ ICML 2024.
- Co-organizer, Deep Learning Theory Workshop & Summer School, Simons Institute, 2022.