SPENCER FREI

frei@berkeley.edu

RESEARCH INTERESTS

Machine learning, statistics, and optimization. Particular interest in theoretical foundations of deep learning, such as neural network optimization, generalization, self-training, and robustness.

ACADEMIC EMPLOYMENT

Postdoctoral Fellow, UC Berkeley and Simons Institute for the Theory of Computing	2021–present
–Mentors: Peter Bartlett and Bin Yu.	1
-Part of the NSF/Simons Collaboration on the Theoretical Foundations of Deep Learning.	
EDUCATION	

 Ph.D, Statistics, UCLA –Advisors: Quanquan Gu and Ying Nian Wu. –Thesis: Statistical learning with neural networks trained by gradient descent. 	2015–2021
MSc., Mathematics, University of British Columbia, Vancouver – Advisor: Ed Perkins. – Thesis: A lower bound for the critical probability in range- <i>R</i> bond percolation.	2013–2015
BSc., Mathematics, McGill University, Montréal –First class honours.	2009–2013

INDUSTRY EXPERIENCE

Research Scientist Intern, Amazon Alexa AI, Cambridge, MA –Worked on natural language understanding using Transformer-based multilingual language mod	<i>Summer 2020</i> lels.
Deep Learning Engineer, Chatterbaby/UCLA, Los Angeles, CA –Used deep learning to detect and analyze infant audio cries recorded in the Chatterbaby app.	2018–2021
Biostatistical Consultant, Ritter Pharmaceuticals, Los Angeles, CA -Consulted on the analysis of clinical trial data using linear regression and mixed effects models.	2017–2019
Statistical Consultant, BlackThorn Therapeutics/UCLA, Los Angeles, CA -Consulted on the analysis of MRI and neuropsychiatric data using penalized regression models.	2016–2018

PUBLICATIONS

Preprints

- 1. **Spencer Frei** and Quanquan Gu. Proxy convexity: a unified framework for the analysis of neural networks trained by gradient descent. Preprint, 2021.
- 2. **Spencer Frei***, Difan Zou*, Zixiang Chen*, and Quanquan Gu. Self-training converts weak learners to strong learners in mixture models. Preprint, 2021.

Refereed Conference Publications

3. Difan Zou*, **Spencer Frei***, and Quanquan Gu. Provable robustness of adversarial training for learning halfspaces with noise. *International Conference on Machine Learning (ICML)*, 2021.

- 4. **Spencer Frei**, Yuan Cao, and Quanquan Gu. Provable generalization of SGD-trained neural networks of any width in the presence of adversarial label noise. *International Conference on Machine Learning* (*ICML*), 2021.
- 5. **Spencer Frei**, Yuan Cao, and Quanquan Gu. Agnostic learning of halfspaces with gradient descent via soft margins. *International Conference on Machine Learning (ICML)*, 2021. **Long talk**.
- 6. **Spencer Frei**, Yuan Cao, and Quanquan Gu. Agnostic learning of a single neuron with gradient descent. *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- 7. Spencer Frei, Yuan Cao, and Quanquan Gu. Algorithm-dependent generalization bounds for overparameterized deep residual networks. *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.

Journal Publications

- 8. Ariana E. Anderson, Mirella Diaz-Santos, **Spencer Frei** *et al.* Hemodynamic latency is associated with reduced intelligence across the lifespan: an fMRI DCM study of aging, cerebrovascular integrity, and cognitive ability. *Brain Structure and Function*, 2020.
- 9. Spencer Frei and Edwin Perkins. A lower bound for p_c in range-R bond percolation in two and three dimensions. *Electronic Journal of Probability* 21(56), 2016.
- 10. **Spencer Frei**, Kathryn Lockwood, Greg Stewart, Justin Boyer, and Burt S. Tilley, On thermal resistance in concentric residential geothermal heat exchangers. *Journal of Engineering Mathematics* 86(1), 2014.

* indicates equal contribution.

AWARDS

• Best Reviewer Award, International Conference on Machine Learning (ICML), 2020 (Top 33%) and 2021 (Top 10%).

- Dissertation Year Fellowship, UCLA Graduate Division, 2020–2021.
- Most Promising Computational Statistician, UCLA Department of Statistics, 2016.
- Research Fellowship, Montréal Institut des sciences mathématiques, 2012.

INVITED TALKS

- "Generalization of SGD-trained neural networks in the presence of adversarial label noise"
- ETH Zürich Young Data Science Researchers Seminar, April 2021.
- Johns Hopkins University Machine Learning Seminar, April 2021.
- Max-Planck-Institute MiS Machine Learning Seminar, March 2021.
- NSF/Simons Collaboration on the Foundations of Deep Learning Seminar, February 2021.

PROFESSIONAL SERVICE

- Reviewer for conferences: ICML 2020 (Top 33% Reviewer), NeurIPS 2020, AISTATS 2021, ICML 2021 (Expert Reviewer; Top 10% Reviewer), NeurIPS 2021, ICLR 2022.
- Reviewer for journals: SIAM Journal on Mathematics of Data Science (SIMODS).
- Reviewer for workshops: Theory of Overparameterized Machine Learning (TOPML) 2021, ICML Workshop on Overparameterization: Pitfalls & Opportunities (ICMLOPPO) 2021.
- Volunteer for Queer in AI.

CONFERENCE, WORKSHOP, AND PROGRAM PARTICIPATION

- International Conference on Machine Learning (ICML). Online, 2021.
- Theory of Overparameterized Machine Learning (TOPML) Workshop. Rice University (online), 2021.
- Neural Information Processing Systems (NeurIPS). Online, 2020.

- Theory of Deep Learning Special Quarter. TTIC/Northwestern Institute for Data, Econometrics, Algorithms, and Learning (online) 2020.
- Neural Information Processing Systems (NeurIPS). Vancouver, BC, 2019.
- Foundations of Deep Learning Program. Simons Institute for the Theory of Computing, 2019.
- Summer School in Probability. Pacific Institute for the Mathematical Sciences, 2014.

TEACHING EXPERIENCE

• UCLA, Department of Statistics

— Guest Lecturer, Stats 231A: Pattern Recognition and Machine Learning, Fall 2020. Lecture on theory of deep learning and the neural tangent kernel (link).

- TA, Stats 100C: Linear Models, Spring 2020.
- TA, Stats 102C: Monte Carlo Methods, Fall 2019.
- TA, Stats 100B: Mathematical Statistics, Winter 2016.
- TA, Stats 10: Intro to Statistics, Summer 2016.
- TA, Stats 100A: Probability Theory, Fall 2016.

PROGRAMMING LANGUAGES

Python, TensorFlow, R, Matlab, AWS (EC2, S3).