

SPENCER FREI

Google Scholar

spencerfrei@ucla.edu

RESEARCH INTERESTS

- Theory of deep learning: optimization, generalization, etc.
- Statistical learning theory
- Applications of deep learning: natural language understanding, audio analysis, etc.

ACADEMIC APPOINTMENTS

- Postdoctoral Fellow, UC Berkeley and Simons Institute for the Theory of Computing *From August 2021*
- Mentors: Peter Bartlett and Bin Yu.
 - Funded by the NSF/Simons program Collaboration on the Theoretical Foundations of Deep Learning.

EDUCATION

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

PROFESSIONAL EXPERIENCE

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

PUBLICATIONS

Refereed Conference Publications

1. Difan Zou*, **Spencer Frei***, and Quanquan Gu. Provable robustness of adversarial training for learning halfspaces with noise. *International Conference on Machine Learning (ICML)*, 2021.
[arXiv:2104.09437](https://arxiv.org/abs/2104.09437).
2. **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.
[arXiv:2101.01152](https://arxiv.org/abs/2101.01152).
Appeared at the Theory of Overparameterized Machine Learning (TOPML) workshop.

3. **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.** [arXiv:2010.00539](https://arxiv.org/abs/2010.00539).
4. **Spencer Frei**, Yuan Cao, and Quanquan Gu. Agnostic learning of a single neuron with gradient descent. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. [arXiv:2005.14426](https://arxiv.org/abs/2005.14426), [conference paper](#).
5. **Spencer Frei**, Yuan Cao, and Quanquan Gu. Algorithm-dependent generalization bounds for overparameterized deep residual networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019. [arXiv:1910.02934](https://arxiv.org/abs/1910.02934), [conference paper](#).

Journal Publications

6. 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. [Journal version](#).
7. **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. [Journal version](#).
8. **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. [Journal version](#).

* indicates equal contribution.

AWARDS

Best Reviewer Award, International Conference on Machine Learning (ICML), 2020.
 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/workshops: ICML 2020 (Best Reviewer), NeurIPS 2020, AISTATS 2021, ICML 2021 (Expert Reviewer), TOPML 2021, NeurIPS 2021.
 Reviewer for journals: SIAM Journal on Mathematics of Data Science (SIMODS).
 Volunteer for Queer in AI.

WORKSHOP AND CONFERENCE PARTICIPATION

Theory of Overparameterized Machine Learning (TOPML) Workshop. Online, 2021.
Neural Information Processing Systems (NeurIPS). Vancouver, BC, 2020.
Theory of Deep Learning Special Quarter. IDEAL (TTIC/Northwestern), Chicago, IL, 2020.
Neural Information Processing Systems (NeurIPS). Vancouver, BC, 2019.
Foundations of Deep Learning. Simons Institute for Computing, Berkeley, CA, 2019.
PIMS Summer School in Probability. Vancouver, BC, 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).