

DANIEL BEAGLEHOLE

Email: dbeaglehole@ucsd.edu Website: dmbeaglehole.github.io

EDUCATION

University of California, San Diego

La Jolla, CA

- Ph.D. in Computer Science and Engineering
Advisor: Mikhail Belkin

Sep 2021 – Present

Columbia University

New York, NY

- M.S. in Computer Science (Foundations Track)
Advisor: Alexandr Andoni

Sep 2019 – Feb 2021

University of Chicago

Chicago, IL

- B.S. in Mathematics with General Honors; Minor in Chemistry

Sep 2015 – Jun 2019

PUBLICATIONS

1. A. Radhakrishnan, **D. Beaglehole***, P. Pandit, M. Belkin. *Mechanism for feature learning in neural networks and backpropagation-free machine learning models*. Science.
2. **D. Beaglehole***, A. Radhakrishnan*, E. Boix-Adserà, M. Belkin. *Toward universal steering and monitoring of AI models*. Under review at Science. Available at arXiv:2502.03708.
3. N. Mallinar, **D. Beaglehole**, L. Zhu, A. Radhakrishnan, P. Pandit, M. Belkin. *Emergence in non-neural models: grokking modular arithmetic via average gradient outer product*. International Conference on Machine Learning (ICML) 2025 [Oral, Top 1% of submitted papers].
4. **D. Beaglehole***, P. Súkeník*, M. Mondelli, M. Belkin. *Average gradient outer product as a mechanism for deep neural collapse*. Conference on Neural Information Processing Systems (NeurIPS) 2024.
5. **D. Beaglehole**, M. Belkin, P. Pandit. *On the inconsistency of kernel ridgeless regression in fixed dimensions*. SIAM Journal on Mathematics of Data Science (SIMODS); also at DeepMath 2022.
6. **D. Beaglehole***, M. Hopkins*, D. Kane*, S. Liu*, S. Lovett*. *Sampling Equilibria: Fast No-Regret Learning in Structured Games*. Symposium on Discrete Algorithms (SODA) 2023.
7. A. Andoni*, **D. Beaglehole***. *Learning to Hash Robustly, Guaranteed*. International Conference on Machine Learning (ICML) 2022.
8. **D. Beaglehole**, I. Mitliagkas, A. Agarwala. *Feature learning as alignment: a structural property of gradient descent in non-linear neural networks*. Transactions on Machine Learning Research (TMLR) 2024.
9. **D. Beaglehole***, A. Radhakrishnan*, P. Pandit, M. Belkin. *Mechanism of feature learning in convolutional neural networks*. Available at arXiv:2309.00570.
10. S. Orsinelli-Rivers, **D. Beaglehole**, T. Einav. *CAPYBARA: A Generalizable Framework for Predicting Serological Measurements Across Human Cohorts*. Available at medRxiv:25331040.
11. T. Valles, **D. Beaglehole**. *Fast, optimal, and dynamic electoral campaign budgeting by a generalized Colonel Blotto game*. Available at arXiv:2406.15714.

INDUSTRY EXPERIENCE

Google DeepMind

Mountain View, CA

- Student Researcher

May 2023 – Mar 2024

Researched theoretical mechanisms for feature learning in neural networks with Atish Agarwala and Ioannis Mitliagkas.

- Summer Research Intern

Jun – Sep 2021

Developed a reinforcement-learning algorithm for asset-price prediction on the Core ML team.

AWARDS

- ARCS Foundation Scholar, UCSD CSE (2023-2025): \$10,000 annual scholarship awarded to a single Ph.D. student in each department
- Summit for AI Institutes Leadership, NSF AI Institute TILOS (2025): selected as the single student to represent UC San Diego.
- Biological Sciences Collegiate Division Summer Fellowship (2016, 2017)
- University of Chicago Dean's List (2016-2018)

INVITED PRESENTATIONS

- Summit for AI Institutes Leadership, NSF AI Institute TILOS (2025): *xRFM: Accurate, scalable, and interpretable feature learning models for tabular data*
- Mathematical and Scientific Foundations of Deep Learning Annual Meeting (2025): *Steering and Monitoring Large Language Models through Feature Learning*
- INFORMS Annual Meeting (2025): *Steering and Monitoring Large Language Models through Feature Learning*
- Simons Foundation (2025): *Steering and Monitoring Large Language Models through Feature Learning*
- UW-Madison ML-OPT Idea Seminar (2023): *Mechanism of feature learning in neural networks*
- Princeton ML Theory Summer School (2023): *Feature learning in neural networks and kernel machines that recursively learn features*
- UCSD Data Science Industry Day (2023): *Feature learning in neural networks and kernel machines that recursively learn features*
- Google Brain (2023): *Feature learning in neural networks and kernel machines that recursively learn features*
- Yale University – Inference, Information, and Decision Systems Group (2023): *Feature learning in neural networks and kernel machines that recursively learn features*
- Symposium on Discrete Algorithms (2023): *Sampling Equilibria: Fast No-Regret Learning in Structured Games*
- International Conference on Machine Learning (2022): *Learning to Hash Robustly, Guaranteed*
- Conference on the Mathematical Theory of Deep Neural Networks (2022): *On the inconsistency of kernel ridgeless regression in fixed dimensions*
- UCSD Theory Seminar (2021): *Learning to Hash Robustly, Guaranteed*
- Goldman Sachs Data Science and Machine Learning Paper Club (2021): *Learning to Hash Robustly, Guaranteed*
- Goldman Sachs Summer Internship Final Presentation (2021): *Predictive Clustering Time Series for Finance*
- Biological Sciences Collegiate Division Fall Symposium (2016, 2017)
- University of Chicago Undergraduate Research Symposium (2016)

ACADEMIC SERVICE

- Journal Reviewer: *Mathematics of Operations Research* (2023); *Neural Computation* (2024)
- Conference Reviewer: SODA 2025, NeurIPS 2025, ICLR 2026, MOSS @ NeurIPS 2025, HiLD @ NeurIPS 2025
- Area Chair: Workshop on the Theoretical Foundations of Foundation Models, ICML 2024