Divit Rawal

divit.rawal@berkeley.edu | Google Scholar | LinkedIn | divitr.github.io

Research Interests: statistical learning theory and learning dynamics; probabilistic and geometric formulations of control and generalization; stochastic processes; statistical physics of deep learning; randomized algorithms.

Education

University of California, Berkeley — Berkeley, CA
B.A. Statistics, Physics, Computer Science — *2023 – 2027*

Relevant Coursework: Statistical Learning Theory†; Theoretical/Mathematical Statistics†; Convex Optimization; Learning for Dynamics and Control†; Probability and Random Processes; Statistical/Thermal Physics; Randomized Linear Algebra, Optimization, Large-Scale Learning†; Statistical Computing. († denotes graduate-level coursework)

Teaching: Head TA (Sp26), Reader (Sp25) for Communication Networks (EE 122). Independently developed lab curriculum, wrote exams, and led weekly recitation sessions.

Publications and Preprints

Minimax Rates for Hyperbolic Hierarchical Learning — *Under review at COLT*, 2026
Divit Rawal, Sriram Vishwanath — [arXiv]

ALPHANSO: Open-Source (α,n) Neutron Source Terms — *Under review at NIM-A*, 2026
Divit Rawal, Anthony J. Nelson, William Zywiec, Daniel Siefman — [GitHub]

Awards and Grants

VESSL AI Academia Grant — Aug 2025

- Academic grant for independent research into embeddings and information geometry; only undergraduate recipient (typically reserved for Ph.D. students and postdoctoral researchers).

Research Experience

Nirvana AI — Jan 2026 – Present
Machine Learning Researcher — *New York, NY*

- Investigating energy-based learning frameworks as alternatives to backpropagation, focusing on implicit differentiation, equilibrium models, and learning dynamics defined by energy minimization.

Lawrence Livermore National Laboratory — Mar 2025 – Present
Researcher (Nuclear Science and Security Consortium Fellow) — *Livermore, CA*

- Developed ALPHANSO, an open-source computational framework for deterministic modeling of (α,n) neutron production using continuous slowing-down approximations and modern evaluated nuclear data.
- Formulated and implemented stochastic and deterministic transport models for neutron yield and spectral prediction, benchmarking against legacy codes and experimental datasets.
- Conducted quantitative validation against experimental neutron yield and spectral benchmarks, analyzing model error and sensitivity to nuclear cross-section data; demonstrated improved predictive accuracy and extensibility relative to legacy SOURCES-based pipelines.

Berkeley Artificial Intelligence Research (BAIR) — Sep 2024 – Present
Researcher (DeWeese Lab) — *Berkeley, CA*

- Analyzed in-context learning in kernel ridge regression via eigenlearning, deriving mode-wise error dynamics and conditions for generalization without representation change.
- Connected ICL behavior to spectral bias in kernel methods (eigenvalue decay, target alignment), clarifying how contextual examples induce effective adaptation in fixed-feature models.
- Ongoing: developing a commutator-based theory of feature acquisition from kernel evolution, predicting stage-wise learning order/timing and diagnosing feature interference via kernel-response dynamics.

UC Irvine, Department of Physics & Astronomy — Feb 2022 – Jul 2023
Researcher (Whiteson Lab) — *Irvine, CA*

- Developed deep learning models using TensorFlow/Keras for high-energy particle collision analysis, achieving 90% accuracy (vs. 80% baseline) in particle identification.
- Simulated particle collisions using MadGraph, Pythia8, Delphes, and ROOT; implemented reconstruction algorithms in C++ and Python achieving 2% mass prediction error.
- Contributed to arXiv:2412.01600 and arXiv:2412.01548, both published in *Reports on Progress in Physics*.

Industry Experience

Cisco Systems (Foundation AI) Jun 2025 – Aug 2025
Software Engineer (Intern) *San Francisco, CA*

- Contributed to post-training and evaluation of cybersecurity-specialized LLMs built on Llama 3.1 8B, including reasoning, instruction tuning, and preference alignment (RLHF) for security dialogue and workflows.
- Supported open-weight releases Foundation-Sec-8B and Foundation-Sec-8B-Instruct, and contributed to the accompanying technical reports documenting data curation, training methodology, and benchmark evaluation.

ExperienceFlow AI May 2024 – Sep 2024
Machine Learning Engineer (Intern) *Remote*

- Modeled finite state machine (FSM) transition dynamics using transformer, state-space, and graph architectures, evaluating generalization and sample efficiency in low-label sequence prediction settings.
- Developed reinforcement-learning-based control policies (Deep Q-Learning, SARSA) for dynamic state reasoning systems, improving stability and convergence in sequential decision tasks.

Amazon Aug 2023 – Dec 2023
Software Engineer (Intern) *Remote*

- Implemented a K-means clustering module in Java for the OpenSearch ml-commons repository, contributing to distributed ML functionality and increasing unit test coverage from 66% to 78%.
- Diagnosed and resolved production data pipeline failures affecting 1M+ users, improving system reliability and reducing query latency in large-scale ML search infrastructure.

Technical Skills

Programming: R, Python, C++, Java, MATLAB, Julia, SQL, Mathematica, LaTeX.
Scientific Computing: NumPy, SciPy, SciKit-Learn, Monte Carlo methods, randomized algorithms.
ML/AI: PyTorch, JAX, TensorFlow, OpenAI Gym, HuggingFace, Isaac Sim.
Tools: Git, Google Cloud, AWS, Docker, Linux, Jupyter, Matplotlib, Pandas, BeautifulSoup, Selenium.