

Gil Goldshlager

ggoldsh@berkeley.edu — ggoldshlager.com

Education

Ph.D. in Applied Mathematics, University of California, Berkeley 2022–Present

Advisor: Lin Lin

B.S. in Mathematics with Computer Science, Massachusetts Institute of Technology 2013–2017

Research Areas

Neural quantum states, scientific machine learning, randomized linear algebra, optimization

Awards

Bernard Friedman Memorial Prize in Applied Mathematics (2026)

Simons Dissertation Fellowship in Mathematics (2025–2027)

Hearts to Humanity Eternal Graduate Research Grant (2025)

Department of Energy Computational Science Graduate Fellowship (2022–2026)

NSF Graduate Research Fellowship Program (2022–2025, declined)

Morais and Rosenblum Undergraduate Research Opportunities Program Award (2014)

Publications

Gil Goldshlager, Jiang Hu, and Lin Lin. A sketch-and-project analysis of subsampled natural gradient algorithms. Forty-third International Conference on Machine Learning (ICML 2026).

Gil Goldshlager, Jiang Hu, and Lin Lin. Worth their weight: randomized and regularized block Kaczmarz algorithms without preprocessing. SIAM Journal on Matrix Analysis and Applications (SIMAX, to appear).

Andres Guzman-Cordero, Felix Dangel, Gil Goldshlager, and Marius Zeinhofer. Improving energy natural gradient descent through Woodbury, momentum, and randomization. Thirty-ninth Annual Conference on Neural Information Processing Systems (NeurIPS 2025).

Ni Zhan, William A. Wheeler, Gil Goldshlager, Elif Ertekin, Ryan P. Adams, and Lucas K. Wagner. Expressivity of determinantal ansatzes for neural network wave functions. Journal of Chemical Theory and Computation (JCTC 2025).

Ethan Epperly, Gil Goldshlager, and Robert J. Webber. Randomized Kaczmarz with tail averaging. Applied and Computational Harmonic Analysis (ACHA 2025).

Nilin Abrahamsen, Zhiyan Ding, Gil Goldshlager, and Lin Lin. Convergence of variational Monte Carlo simulation and scale-invariant pre-training. Journal of Computational Physics (JCP 2024).

Gil Goldshlager, Nilin Abrahamsen, and Lin Lin. A Kaczmarz-inspired approach to accelerate the optimization of neural network wavefunctions. *Journal of Computational Physics* (JCP 2024).

Jeffmin Lin, Gil Goldshlager, and Lin Lin. Explicitly antisymmetrized neural network layers for variational Monte Carlo simulation. *Journal of Computational Physics* (JCP 2023).

Royce Rozov, Gil Goldshlager, Eran Halperin, and Ron Shamir. Faucet: streaming de novo assembly graph construction. *Bioinformatics* (2018).

Invited Talks

A sketch-and-project approach to natural gradients in scientific machine learning.
Workshop on Geometric Methods in Scientific Machine Learning, Oberwolfach, August 2026.

Towards high-precision optimizers for neural wavefunctions.
Electronic Structure Workshop, University of Wisconsin–Madison, June 2026.

Towards high-precision optimizers for neural wavefunctions.
Total Energy Mini-Workshop, UC Berkeley, January 2026.

Understanding and accelerating subsampled natural gradient descent.
Colloquium in Applied and Computational Mathematics, ETH Zurich, September 2025.

Accelerating stochastic reconfiguration in the under-sampled regime.
Workshop on Stochastic Electronic Structure Methods, Telluride, June 2025.

Contributed Talks

Towards high-precision optimizers for scientific machine learning.
Conference of the International Linear Algebra Society, Virginia Tech, May 2026.

Elucidating subsampled optimizers for variational Monte Carlo simulations.
APS Global Physics Summit, Denver, March 2026.

Explaining the success of subsampled natural gradient algorithms in scientific machine learning.
European Conference on Numerical Mathematics and Advanced Applications, Heidelberg, September 2025.

Seminar Talks

New insights on subsampled natural gradient methods for scientific machine learning.
Michael Mahoney Group, UC Berkeley, October 2025.

A randomized linear algebra perspective on subsampled natural gradient methods.
Recent Progress and Open Directions in Matrix Computations, Simons Institute for the Theory of Computation, September 2025.

Understanding and accelerating subsampled natural gradient descent.
Juan Carrasquilla-Alvarez Group, ETH Zurich, September 2025.

Understanding and accelerating subsampled natural gradient descent.
Computational Quantum Science Laboratory, EPFL, September 2025.

Using neural network wavefunctions to advance electronic structure theory.
Bakar Institute of Digital Materials for the Planet, UC Berkeley, September 2024.

SPRING: a new optimizer to accelerate the training of neural network wavefunctions.
Yuehaw Khoo Group, University of Chicago (virtual), February 2024.

Other Activities

Referee for: SIAM Journal on Matrix Analysis and Applications, SIAM Journal on Scientific Computing, Applied Mathematics and Computation, Inverse Problems, International Conference on Machine Learning, Machine Learning: Science and Technology, Journal of Computational Physics.

Undergraduate research mentor for: Jeffery Ren (SULI program, June 2025–August 2025; continued research, September 2025–present), Brij Patel (April 2026–present).

Teaching assistant for Advanced Algorithms, Massachusetts Institute of Technology, Fall 2014.

Work Experience

Research intern, Lawrence Berkeley National Laboratory	Summer 2025
Supervisor: Bert de Jong	
Consultant, Actuate Innovation	2020–2021
Software engineer, Coda Project Inc.	2017–2019
Bioinformatics intern, 20n Labs	June–December 2016
Software engineering intern, Google	Summer 2014