Corey Adams, Ph.D.

Senior Software Engineer | Computational Scientist | Physicist

LinkedIn • GitHub • Google Scholar

Interests: Deep Learning, High Performance Computing, AI for Physics

Background: Accomplished computational scientist with a strong foundation in fundamental science and a track record of applying cutting-edge AI methodologies to diverse projects. A decade of experience with performance optimization for AI and HPC, from single device to supercomputer scale. Proficient in integrating AI solutions into complex computing environments, pushing the boundaries of high-performance computing and contributing to the intersection of AI, machine learning, and computational science.

Skills

Deep Learning: Computer Vision (CNN, GNN, Classification, Detection, Segmentation, UNet, YOLO), Sparse Neural Networks, Generative, Unsupervised and Self-Supervised Learning (GANs, AEs, SimCLR), Surrogate Models, Neural Operators, Distributed Learning and Inference, Scaling, Graph Neural Networks, Machine Learning, Pattern Recognition.

Languages and Frameworks: PyTorch (including FSDP, DTensor), Python, C++, C, CUDA-Python, TensorFlow, Keras, JAX.

Distributed & HPC: MPI, NCCL/RCCL/oneCCL, Slurm/PBS, HDF5/Zarr, Docker/Enroot.

Profiling and Performance: Nsight Systems, VTune, PyTorch Profiler, cProfile/line_profiler (python), mixed precision, CUDA graphs, Triton, NVIDIA Warp, torch.compile.

Foundational: Mathematics (differential equations, linear algebra, advanced calculus), data analysis, probability and statistics, fundamental physics, mathematical modeling, research skills.

Communication: Technical and popular writing; creator of lectures and workshops on AI and scientific computing; management and mentorship of staff; teamwork and collaboration.

Experience

NVIDIA Corporation Dec 2024 – Present

Senior Software Engineer, PhysicsNeMo

- **Domain Parallelism:** Invented and architected ShardTensor, a novel domain-parallelism framework that scales high-resolution scientific data beyond single-GPU memory limits during training and inference. Now a core capability of PhysicsNeMo. Documentation, Examples, Publication.
- **Performance:** Served as performance lead for PhysicsNeMo, delivering up to 30x speedups on top CAE models (Blog post). Architected SciML datapipes for high-concurrency IO with GPU stream parallelism.
- **GeoTransolver:** Authored GeoTransolver (jointly with R. Ranade), PhysicsNeMo's top-performing surrogate model for CFD and CAE applications. Training Recipe, Publication.
- **v2.0 Leadership:** Led the entire engineering team in the v2.0 refactor of PhysicsNeMo, delivering on time and leading to a sustained increase in downloads. Monthly downloads increased by 3x since the v2.0 release.

- Top contributor to PhysicsNeMo (3k+ GitHub stars, 100+ contributors) over the past 12 months, owning domain parallelism, datapipes, multiple models, CI/CD, testing infrastructure, and multiple end-to-end training recipes.

Advanced Micro Devices (AMD) July 2024 – Dec 2024
Senior Member of Technical Staff

- Lead for AI on the ROCm stack and Instinct GPU product line for Technical Marketing Engineering.
- Developed content, demos, and technical and educational material related to AI GPU programming on MI300X GPUs.

Argonne National Laboratory Sept 2018 – July 2024
Computational Scientist, Physicist (Oct 2022 – July 2024); Asst. Computer Scientist, Asst. Physicist (Sept 2018 – Oct 2022)

- Lead of AI software on the Aurora supercomputer, #2 in the world with more than 60k GPUs.
- Developed multiple novel AI techniques for science, including graph neural networks that obey fundamental physical laws in JAX to solve quantum many-body physics.
- Created a unique, fully differentiable simulation software in JAX where trainable parameters are based on interpretable physical quantities of experimental devices, trained with both supervised and unsupervised methods.
- Interfaced with hardware vendors for benchmarking AI-for-science applications, including NVIDIA, SambaNova, and Cerebras.

Harvard University Feb 2017 – Sept 2018
Postdoctoral Research Associate

- Physicist working on the MicroBooNE, SBND, and NEXT neutrino experiments.
- Pioneered applications of deep learning to neutrino physics.

Education

- Ph.D., Physics — Yale University May 2017
- M.Phil., Physics — Yale University May 2014
- M.Sci., Physics — Yale University May 2014
- B.S., Mathematics — University of Rochester May 2011
- B.S., Physics — University of Rochester May 2011

Sample Projects

- **PhysicsNeMo** — Open-source AI physics software from NVIDIA. 3000+ stars on GitHub. Top contributor since joining the team in Dec. 2024, of 100+ contributors. Lead architect of domain parallelism, the datapipes, and multiple training recipes, as well as the core architect for the repository.
- **CosmicTagger** — Deep-learning image segmentation (UResNet) for fundamental physics detectors. Beats state-of-the-art by 5x accuracy; deployed as a comparative benchmark and acceptance test on multiple supercomputers. PyTorch, JAX, and TF implementations. Scaled to thousands of ranks on multiple HPC systems. (Lead Developer, Inactive)
- **Larcv3** — C/C++ data serialization package for sparse and irregular data, built on HDF5 with a Python interface. Scaled to 20k+ ranks with MPI. (Lead Developer, Inactive)
- **SparseEventID** — Sparse, 3D computer vision models for fundamental physics detectors. Classification as well as object detection and SimCLR unsupervised representation learning. PyTorch. (Lead Developer, Inactive)

- **QMC** — Differentiable solver of the multibody Schrödinger equation (quantum mechanics) in JAX. Scales with mpi4jax to thousands of ranks. (Lead Developer, Inactive)

Awards

- Department of Energy Early Career Award ($2.5M) — Office of Science, Nuclear Physics 2021
- Impact Argonne Awards (4) — Argonne National Laboratory 2018–2022
- Neutrino Physics Center Scholar — Fermi National Accelerator Laboratory 2016
- University Research Association Visiting Scholar — Fermi National Accelerator Laboratory 2015
- Excellence in Undergraduate Teaching Award — University of Rochester 2011

Selected Publications

Google Scholar (70+ more, h-index 40)

- "GeoTransolver: Learning Physics on Irregular Domains Using Multi-scale Geometry Aware Physics Attention Transformer," C. Adams, R. Ranade, R. Cherukuri, S. Choudhry, Submitted to NeurIPS 2026.
- "ShardTensor: Domain Parallelism for Scientific Machine Learning," C. Adams, P. Harrington, A. Subramaniam, M. S. Abbas, J. Pathak, M. Pritchard, S. Choudhry, Submitted to SuperComputing 2026.
- "Demonstration of Portable Performance of Scientific Machine Learning on High Performance Computing Systems," K. Hossain, R. Balin, C. Adams, et al., Proceedings of the SC'23 Workshops (2023).
- "Nuclei with up to A = 6 nucleons with artificial neural network wave functions," A. Gnech, C. Adams, N. Brawand, G. Carleo, A. Lovato, N. Rocco, Few-Body Systems, 63, 1 (2022).
- "Variational Monte Carlo calculations of $A \leq 4$ nuclei with an artificial neural-network correlator ansatz," C. Adams, G. Carleo, A. Lovato, N. Rocco, Phys. Rev. Lett. 127, 022502 (2021).
- "Demonstration of background rejection using deep convolutional neural networks in the NEXT experiment," NEXT Collaboration, Journal of High Energy Physics, vol. 2021, art. 189 (2021).
- "Cosmic Background Removal with Deep Neural Networks in SBND," SBND Collaboration, Front. Artif. Intell. (24 August 2021).
- "Convolutional Neural Networks Applied to Neutrino Events in a Liquid Argon Time Projection Chamber," MicroBooNE Collaboration, JINST 12, P03011 (2017).

Selected Workshops, Lectures and Seminars

- AI Frameworks on the Aurora Supercomputer — ALCF Developer Session, Virtual, 2024
- AI for Fundamental Physics — DOE National Science Bowl, Science Day, Leesburg, VA, 2024
- The Aurora Supercomputer — Argonne Colloquium, Lemont, IL, 2024
- Neutrinoless Double Beta Decay at Argonne — Argonne Seminar, Lemont, IL, 2024
- CNNs and GNNs for Physics — COFI Winter School, San Juan, PR, 2023
- AI in the NEXT Experiment — AI in Nuclear Physics, DNP 2023
- Calculation of Nuclear Ground States using Artificial Neural Networks — ALCF Hands-on Workshop, Lemont, IL, 2023
- AI and Neutrino Physics — U. Chicago Physics Seminar, Hyde Park, IL, 2023
- Sparse IO for Deep Learning — PyHEP Workshop, Virtual, 2022
- Machine Learning Quantum Mechanics — Argonne Training Program for Extreme Scale Computing, St. Charles, IL, 2022
- AI/ML on High Performance Computing — Beam Dynamics Workshop, Chicago, IL, 2022
- AI in Neutrinoless Double Beta Decay — Physics Seminar, Rice University, 2022
- AI and Neutrinoless Double Beta Decay — MODE Conference, Virtual, 2021

- Artificial Intelligence – Computing Requirements — National Nuclear Security Administration, Emerging Technologies, Virtual, 2021
- Machine Learning for Neutrinoless Double Beta Decay in NEXT — Dark-matter and Neutrino Computation Explored (DANCE) ML, Virtual, 2020
- Machine Learning Day, "Autoencoders" — Argonne Training Program for Extreme Scale Computing, Virtual, 2020
- Python on HPC — Simulation, Data, and Learning Workshop, Virtual, 2020
- ALCF Data and Learning Frameworks — Computational Performance Workshop, Lemont, IL, 2019
- Sparse Convolutional Networks — Greater Chicago Area Systems Research Workshop, Chicago, IL, 2019
- Machine Learning, Data Science and Neutrino Physics at Argonne's Leadership Computing Facility — SLAC AI Seminar, Menlo Park, CA, 2019
- Deep Learning and Neutrino Physics — Argonne Leadership Computing Facility Seminar, Lemont, IL, 2018

Professional Service

Committee Service: Argonne Named Fellow Committee (2019, member; 2020–2023, co-chair); ASRC Integrated Research Infrastructure Task Force (2020–2021); ALCF Postdoc Committee (chair, 2020–2024); Argonne Postdoc Award Committee (member, 2020); Physics DEIA Committee (member, 2019–2024).

Workshops Led or Co-Organized: ALCF AI Training Series (2022, 2023); Computational Performance Workshop (2021); Neutrino Physics and Machine Learning (2020); Simulation, Data, Learning Workshop (2019).