Nikoli Dryden

ndryden.com
dryden1@llnl.gov • ndryden@ndryden.com

I work at the intersection of deep learning and high-performance computing. I build systems for training deep neural networks fast and at scale, design more efficient neural networks, and apply deep learning to scientific applications.

EDUCATION

2014–2019	**PhD Computer Science**	**University of Illinois at Urbana-Champaign**

Committee: Marc Snir (advisor), William Gropp, Wen-mei Hwu, Alexander Schwing, Brian Van Essen
Thesis: *Large-Scale Training of Deep Neural Networks*

2010–2014	**BS Computer Science**	**University of Illinois at Urbana-Champaign**

Minor in Mathematics, James Scholar

PROFESSIONAL EXPERIENCE

2023–present — Computer Scientist — Lawrence Livermore National Laboratory
Leading research on exascale high-performance & scalable deep learning, and scientific machine learning.

2019–2022 — ETH Postdoctoral Fellow — ETH Zürich
Research on systems for deep learning, efficient neural networks, and deep learning for weather & climate in Torsten Hoefler's Scalable Parallel Computing Laboratory. Mentored ten bachelors and ten masters students.

2014–2019 — Research Assistant — University of Illinois at Urbana-Champaign
Research on algorithms and systems for scalable training of deep neural networks, and runtime systems for future exascale supercomputers in Marc Snir's High-Performance Computing group.

Summer 2016, 2017, 2018 — Computation Intern — Lawrence Livermore National Laboratory
Research on scalable deep learning on supercomputers. Led development of the Aluminum communication library and contributed to the LBANN toolkit for training deep neural networks.

Summer 2015 — Summer Intern — Raytheon Centers of Innovation
Developed tools for cloud-based analysis of OS images and large-scale indexing/search of unstructured data.

NOTABLE PROJECTS

2024–present — ElMerFold
Technical lead for a scalable and high-performance bulk inference framework. Produced the OpenFold 3 protein structure distillation datasets.

2024–present — LBANNv2 [github.com/LBANN]
Reimagining of the LBANN framework as a PyTorch backend and modular supporting libraries. Major contributions to distributed training algorithms and AMD MI300A-specific optimizations.

2019–2022 — DaCeML [github.com/spcl/daceml & github.com/spcl/substation]
Co-led research and development of a deep learning compiler based on data movement analysis. As part of this we created the fastest single-GPU implementation for training BERT.

2016–2024 — LBANNv1 [github.com/LBANN]
Research and development "across the stack" for the Livermore Big Artificial Neural Network toolkit, a homegrown high-performance deep learning training framework. Major contributions to communication infrastructure, GPU acceleration, model-parallelism, and its Python interface.

AWARDS & HONORS

2025 — Director's Exemplary Team Science Award — Lawrence Livermore National Laboratory
For the "exceptional execution in the acquisition, partnership, deployment, and application preparation for the El Capitan Exascale system".

2023 — R&D 100 Award — Lawrence Livermore National Laboratory
Received as part of the Cancer Distributed Learning Environment (CANDLE) team.

2019 — ETH Postdoctoral Fellow — ETH Zürich
Competitive fellowship for "young researchers who have already demonstrated scientific excellence".

2019 — Kenichi Miura Award — University of Illinois at Urbana-Champaign
"Honors a graduate student for excellence in high-performance computing".

2018 — State Farm Doctoral Scholar — University of Illinois at Urbana-Champaign
Fellowship "... to support outstanding continuing doctoral students".

2013 — Student of the Year — National Center for Supercomputing Applications
Awarded once per year to the most outstanding student working at NCSA.