Nikoli Dryden

ndryden.com
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, 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.

2012–2014 **SPIN Fellow** — **National Center for Supercomputing Applications**
Member of the innaugural Students Pushing Innovation (SPIN) program. Subsequently co-mentored a student in the program. Research and development for parallel debugging tools on supercomputers.

Summer 2012 **High Energy Density Physics Intern** — **Lawrence Livermore National Laboratory**
Developed PGDB, a parallel debugger for MPI applications at scale.

Notable Projects

2019–present **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.

2019–2022 **Deep Weather / MAELSTROM** [`github.com/spcl/deep-weather` & `github.com/spcl/ens10`]
Co-led research in deep learning for weather and ETH's role in the EuroHPC MAELSTROM project. Set new state-of-the-arts for neural ensemble post-processing and weather prediction, and released the ENS-10 dataset.

2017–present **Aluminum** [`github.com/LLNL/Aluminum`]
A high-performance GPU communication library. Lead developer.

2016–present **LBANN** [`github.com/LLNL/LBANN`]
Research and development "across the stack" for the Livermore Big Artificial Neural Network toolkit. Major contributions to communication infra, GPU acceleration, model-parallelism, and its Python interface.

Awards & Honors

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.