

# Nikoli Dryden

ndryden@ndryden.com — dryden2@illinois.edu — ndryden.com

## RESEARCH INTERESTS

Scalable training of deep neural networks; applying deep learning to scientific and computational simulation applications; high-performance and parallel computing; irregular applications

## EDUCATION

**BS Computer Science** 2010 - 2014

University of Illinois at Urbana-Champaign  
*Minor in Mathematics, James Scholar*

**PhD Computer Science** 2014 - 2019

University of Illinois at Urbana-Champaign  
Advisor: Prof. Marc Snir  
Thesis: *Large-Scale Training of Deep Neural Networks*

## PROFESSIONAL ETH Zurich

Zurich, Switzerland  
*ETH Postdoctoral Fellow* August, 2019 - present  
Postdoc with Prof. Torsten Hoefer. Large-scale deep learning and its applications.

## Lawrence Livermore National Laboratory

Livermore, CA  
*Computation Intern* May - August, 2016, 2017, 2018  
Developed algorithms and software for large-scale training of deep neural networks. Communication optimization and quantization, distributed-memory convolution, applications of deep learning to scientific and computational simulation datasets. Helped develop the LBANN toolkit.

## University of Illinois

Urbana, IL  
*Research Assistant* May, 2014 - present  
RA for Prof. Marc Snir, researching scalable deep learning and scalable runtimes for exascale supercomputers.  
Further developed scalable deep learning algorithms in collaboration with LLNL. Developed the PPL C++11 parallel runtime and associated applications.

## Raytheon Centers of Innovation

Arlington, VA  
*Intern* May, 2015 - August, 2015  
Cloud-based analyses. Large-scale indexing and search of unstructured data with Python and AWS EMR.

## Lawrence Livermore National Laboratory

Livermore, CA  
*HEDP Intern; Offsite Collaborator* May, 2012 - May, 2015  
Developed an open-source parallel debugger for large-scale MPI applications, PGDB.  
Used PGDB to support the development of Kull, a large multi-physics application.

## NCSA/XSEDE

Urbana, IL  
*SPIN Fellow, Advanced SPIN Fellow* January, 2013 - May, 2014  
Improved scalability and platform support for PGDB.  
Mentored an undergraduate SPIN Fellow working on PGDB.

## KRR Group

Urbana, IL  
*Undergraduate Researcher* February, 2012 - January 2013  
Graphical Kripke models.

**PUBLICATIONS** Yosuke Oyama, Naoya Maruyama, **Nikoli Dryden**, Peter Harrington, Jan Balewski, Satoshi Matsuoka, Marc Snir, Peter Nugent, and Brian Van Essen. “The Case for Strong Scaling in Deep Learning: A Case Study with a 3D Cosmological Model.” In prep.

Alex Brooks, **Nikoli Dryden**, Hoang-Vu Dang, Andrew Siegel, and Marc Snir. “Simple and Fast Monte Carlo Particle Transport Algorithms.” In prep.

Chen Wang, **Nikoli Dryden**, Franck Cappello, and Marc Snir. “Delayed Detection of Silent Errors.” In prep.

**Nikoli Dryden**, Naoya Maruyama, Tim Moon, Tom Benson, Marc Snir, and Brian Van Essen. “Channel and Filter Parallelism for Large-Scale CNN Training.” To appear at *Supercomputing 2019*.

**Nikoli Dryden**, Naoya Maruyama, Tom Benson, Tim Moon, Marc Snir, and Brian Van Essen. “Improving Strong-Scaling of CNN Training by Exploiting Finer-Grained Parallelism.” *IPDPS 2019*.

**Nikoli Dryden**, Naoya Maruyama, Tim Moon, Tom Benson, Andy Yoo, Marc Snir, and Brian Van Essen. “Aluminum: An Asynchronous, GPU-aware Communication Library Optimized for Large-Scale Training of Deep Neural Networks on HPC Systems.” *MLHPC 2018*.

Chen Wang, **Nikoli Dryden**, Franck Cappello, and Marc Snir. “Neural Network Based Silent Error Detector.” *Cluster 2018. (Best Paper)*

Roshan Dathathri, Gurbinder Gill, Loc Hoang, Hoang-Vu Dang, Alex Brooks, **Nikoli Dryden**, Andrew Lenharth, Marc Snir, and Keshav Pingali. “Gluon: A Communication Optimizing Framework for Distributed Heterogeneous Graph Analytics.” *PLDI 2018*.

Hoang-Vu Dang, Roshan Dathathri, Gurbinder Gill, Alex Brooks, **Nikoli Dryden**, Andrew Lenharth, Loc Hoang, Keshav Pingali, and Marc Snir. “A Lightweight Communication Runtime for Distributed Graph Analytics.” *IPDPS 2018*.

Sam Adé Jacobs, **Nikoli Dryden**, Roger Pearce, and Brian Van Essen. “Towards Scalable Parallel Training of Deep Neural Networks.” *MLHPC 2017*.

**Nikoli Dryden**, Tim Moon, Sam Adé Jacobs, and Brian Van Essen. “Communication Quantization for Data-parallel Training of Deep Neural Networks.” *MLHPC 2016*.

Alex Brooks, Hoang-Vu Dang, **Nikoli Dryden**, and Marc Snir. “PPL: An abstract runtime system for hybrid parallel programming.” *ESPM2 2015*.

**Nikoli Dryden**. “PGDB: A Debugger for MPI Applications.” *XSEDE 2014*.

## POSTERS

Naoya Maruyama, **Nikoli Dryden**, Tom Benson, Tim Moon, Brian Van Essen, and Marc Snir. “Spatial Parallel Convolution for Scalable Training of Large Neural Networks.” GPU Technology Conference, 2019.

**Nikoli Dryden**, Naoya Maruyama, Tom Benson, Tim Moon, Marc Snir, and Brian Van Essen. “Scalable CNN Training on Large-Scale HPC Systems.” Workshop on Systems for ML and Open Source Software at NIPS, 2018.

Naoya Maruyama, **Nikoli Dryden**, Tim Moon, Brian Van Essen, and Marc Snir. “Generalized Distributed-Memory Convolutional Neural Networks for Large-Scale Parallel Systems.” LLNL Data Science Institute Workshop, 2018.

Sam Adé Jacobs, **Nikoli Dryden**, Tim Moon, Brian Van Essen, Stewart He, and Jonathan Allen. “Scaling Deep Learning for Cancer Drug Discovery on HPC Systems.” LLNL Data Science Institute Workshop, 2018.

**Nikoli Dryden**, Brian Van Essen, and Marc Snir. “Gradient Quantization for Data-Parallel DNN Training.” Salishan Conference on High Speed Computing, 2017.

## TALKS

**ESPT 2015**

November, 2015

Work-in-progress talk: *Large-scale debugging with graphs*

	<b>NCSA</b> <i>Towards an Improved PGDB</i>	November, 2013
	<b>Argonne National Laboratory</b> <i>A GDB-based Debugger for MPI Applications</i>	October, 2013
	<b>NCSA</b> <i>A Parallelized GDB-based Debugger</i>	April, 2013
	<b>Lawrence Livermore National Laboratory</b> <i>PGDB</i>	August, 2012
<b>AWARDS</b>	<b>ETH Postdoctoral Fellowship</b> From: ETH Zurich	Fall, 2019
	<b>Kenichi Miura Award</b> From: University of Illinois at Urbana-Champaign	Spring, 2019
	<b>Graduate College Travel Grant</b> From: University of Illinois at Urbana-Champaign	Spring, 2019
	<b>State Farm Doctoral Scholar</b> From: University of Illinois at Urbana-Champaign and State Farm	Spring, 2018
	<b>Student of the Year 2013</b> From: National Center for Supercomputing Applications	November, 2013
	<b>“Defining the Future” 1<sup>st</sup> Place</b> From: University of Illinois at Urbana-Champaign Engineering Open House	March, 2011
<b>TEACHING</b>	<b>UIUC CS 420</b> Teaching assistant for Parallel Programming: Science and Engineering	Fall, 2017
<b>SERVICE</b>	<b>Principles and Practice of Parallel Programming</b> Artifact evaluation committee	2019
	<b>Machine Learning in HPC Environments</b> Program committee	2018
<b>PROJECTS/ INVOLVEMENT</b>	<b>LBANN</b> Open-source software framework for scaling the training of deep neural networks on HPC systems. Research involves communication optimization and quantization, scaling data-parallel training, model-parallel training of very large networks, performance modeling, and development of more scalable training algorithms while maintaining model accuracy and generalization.	May, 2016 - present
	<b>Aluminum</b> Open-source communication library providing a generic interface to different communication backends and custom, optimized implementations. A focus is on GPU-centric communication that avoids unnecessary synchronizations. Used in LBANN and the Hydrogen distributed linear algebra library.	January, 2018 - present
	<b>PPL</b> Experimental C++11 parallel runtime for exploring implementation tradeoffs in the context of future exascale supercomputers. Developed proxy applications, including energy band Monte Carlo, Barnes-Hut, and breadth-first search.	May, 2014 - present
	<b>PGDB</b> Open-source debugger for MPI applications at scale, implemented in Python and C. Collaborated with Prof. Marc Snir to further scalability as part of my senior thesis.	May, 2012 - present

**Association for Computing Machinery**

March, 2013 - present

Treasurer of student chapter of the ACM at UIUC, 2013-2014.

Managed funds for the chapter and the Reflections | Projections 2013 through 2016 conferences.

**SIGArt**

September, 2010 - May, 2015

Artificial intelligence projects. Chair from spring 2011 to 2014.

- 2010 - 2011: Multi-touch gesture recognition using hidden Markov models.  
Won first place in the “Theme: Defining the Future” category of Engineering Open House 2011.
- 2011 - 2012: Improved multi-touch techniques through learning HMMs.
- 2012 - 2013: Monocular vision-based realtime obstacle detection for quadrotors.
- 2013 - 2014: Landing zone detection for quadrotors.
- 2014 - 2015: Audio-based gesture recognition and control for quadrotors.

**Operating System**

August, 2011 - December, 2011

Group project developing a concurrent x86-based kernel.

**SKILLS**

<b>Languages</b>	C/C++, Python, Java, MATLAB, L <sup>A</sup> T <sub>E</sub> X, x86, Common Lisp, Lua, others
<b>Frameworks/APIs</b>	MPI, POSIX, CUDA, LBANN, PyTorch, cuDNN, OpenMP, qthreads, others
<b>Tools/Technologies</b>	Emacs, GDB, Valgrind, Vtune/nvprof, LaunchMON, MRNet, others