

Education

Georgia Institute of Technology, M.S. Computer Science (Atlanta, GA) *2017 – 2019*

- » Specialization in High Performance Computing.
- » Research in optimization & machine learning theory, advised by Dr. Jacob Abernethy.

University of Michigan, B.S. Honors Applied Mathematics (Ann Arbor, MI) *2013 – 2017*

Research

Research Intern, RIKEN AIP, advised by Dr. Emtiyaz Khan (Tokyo, Japan) *Summer 2019*

- » Studied training dynamics of variational autoencoders (VAEs) and the amortization gap.
- » Ran experiments in PyTorch; became interested in automatic differentiation and compiler tools for ML.

Graduate Researcher, advised by Dr. Jacob Abernethy (GT) *2017 – Current*

- » Continuous-time analysis of gradient descent and its relatives; study of discretization error.
- » Optimization and online learning on manifolds; natural gradient and projection-free methods.

Digital Humanities & German Periodicals, advised by Dr. Peter McIsaac (UM) *2013 – 2015*

- » Built a flask webapp enabling humanities researchers to run and visualize statistical topic models on a large corpus of 19th-century German periodicals.
- » Corrected noisy digital scans using a Hidden Markov Model over word fragments. Trained on a synthetic dataset of eroded digital scans with known text. Segmented words based on letter-successor-variety.
- » Implemented online variational inference for LDA/HDP from scratch in Python.
- » Trained a logistic regression classifier to detect toponyms using `word2vec` features, achieving near state-of-the-art performance (0.85 F_1 -score) on a Dutch dataset.

Internships

Software Engineering Intern, Microsoft (Seattle, WA) *Summer 2016*

- » Built a multiplatform mobile app with Xamarin to display Windows telemetry insights to developers.

Data Science Intern, Are You a Human (Detroit, MI) *Summer 2015*

- » Designed new features to improve bot classification accuracy by 4%
- » Implemented random forest models for user fingerprinting based on device capabilities.

Teaching

Teaching Assistant, CS 4540, Advanced Algorithms for Machine Learning (GT) *F18, F19*

- » Authored lecture notes, homework, and demonstrations for a flipped-classroom course.
- » Designed a new curriculum with the goal of exposing undergrads to the math behind modern ML.
- » Topics included convex geometry, numerical methods, linear programming, and statistical inference.

Teaching Assistant, EECS 545/445, Machine Learning (UM) *S16, F17, S17*

- » Redesigned the curriculum with Prof. Jacob Abernethy, with emphasis on statistical methods.
- » Taught a weekly discussion section of around twenty students.

Selected Projects

Incompressible Fluid Simulation *2019*

- » Interactive CUDA simulation of incompressible flow using parallel Jacobi solver with vorticity confinement.
- » Compared against CPU implementation of incomplete Cholesky preconditioned conjugate gradient.

Matey, a numerical linear algebra library for Python, written in C++ *2017*

- » Fast matrix operations, factorization, linear system solving, and eigenvalue computations from scratch.
- » Built as a Python C-extension, using CUDA to parallelize existing algorithms.

Technical Experience	<i>Advanced</i>	<i>Proficient</i>	<i>Familiar</i>
Programming Languages	Python	JavaScript/TypeScript, C#, CUDA	C++, Haskell, Go
Machine Learning	numpy, gensim	matplotlib, scikit-learn	pytorch
Miscellaneous		L ^A T _E X, Git, LLVM, flask	AWS, Node, Docker

Relevant Coursework (*audited)

Cs 6290, Advanced Computer Architecture	(Prof. Tom Conte, GT)	<i>S19</i>
» Pipelining; instruction-level parallelism; superscalar processors; VLIW; Tomasulo/ROB/RAT		
» Memory hierarchies; multiprocessors; shared memory vs. message passing; cache coherency/consistency		
Cs 6241, Advanced Compiler Optimizations	(Prof. Santosh Pande, GT)	<i>S19</i>
» For projects, wrote LLVM transform passes to perform optimizations discussed in class.		
» Reaching definitions; available expressions; partial redundancy elimination; SCCP; infeasible paths		
» Data/loop/control dependencies; loop parallelism, reordering, unrolling, and vectorization		
CSE 6220, High-Performance Parallel Computing	(Profs. Aluru & Catalyurek, GT)	<i>S19</i>
» Parallel runtime analysis; efficiency; interconnection networks & embeddings; MPI programming		
» Prefix sum; bitonic sort; sample sort; Cannon's algorithm; parallel FFT		
MATH 7244, Stochastic Calculus	(Prof. Michael Damron, GT)	<i>F18</i>
» Brownian motion; mean-square calculus; continuous-time martingales; stochastic integration		
ISYE 7687, Discrete Optimization for Machine Learning	(Prof. Sebastian Pokutta, GT)	<i>S18</i>
» Boosting and online learning; bandits and reinforcement learning; away-step conditional gradient		
» Completed a final report surveying algorithms for online submodular maximization.		
STATS 700, Bayesian Nonparametrics Seminar	(Prof. Long Nguyen, UM)	<i>F14, S16</i>
» Existence and construction of Dirichlet processes; stick-breaking and Chinese restaurant processes		
» Indian Buffet Processes; Hierarchical Dirichlet Processes; online variational inference and natural gradients		
» Likelihood and sufficiency principles; Bayesian vs. frequentist statistics		
» Exponential families and conjugacy; Gibbs sampling; variational inference		
Cs 7545, Statistical Learning Theory	(Prof. Jacob Abernethy, GT)	<i>F17</i>
» PAC-learning and VC-dimension; margin learning and kernel methods; boosting		
» Online convex optimization; convex-concave games; exponential weights; mirror descent		
Cs 6550, Design & Analysis of Algorithms	(Prof. Jamie Morgenstern, GT)	<i>S18</i>
» Matroids and greedy algorithms; graph connectivity and shortest paths; matchings; linear programming		
» Gradient and mirror descent; ellipsoid method; Johnson-Lindenstrauss and random projections		
» Completed a final report on random matrix theory and algorithms for sampling random matrices.		
MATH 6455, Differential Geometry*	(Prof. Mohammad Ghomi, GT)	<i>S18</i>
» Smooth manifolds; vector fields; geodesics; Riemannian metrics; Levi-Civita connection		
» Submanifolds; differential forms; Lie groups; integral curves and flows		
MATH 571, Numerical Linear Algebra	(UM)	<i>F15</i>
MATH 671, Fast Numerical Methods	(UM)	<i>S17</i>
CSE 8803, Advanced Scientific Computing	(Prof. Edmond Chow, GT)	<i>S18</i>
» Stationary iterative methods; conjugate gradient and Krylov subspaces; Chebychev polynomials		
» Nonuniform FFT and butterfly algorithms; Ewald summation; multigrid; fast multipole methods		
» Finite element analysis; molecular simulation with hydrodynamic interactions		
MATH 597, Measure Theory & Real Analysis	(UM)	<i>S16</i>
MATH 525, Probability Theory	(UM)	<i>F16</i>
MATH 420, Advanced Linear Algebra	(UM)	<i>F15</i>

Other Involvement

President, Michigan Student Artificial Intelligence Lab *2015-2017*

- » Organized a weekly machine learning reading group for undergraduate & graduate students

(Modified: 7 February 2020)