Nand Dalal

📞 +1 (214) 998 5160
✉ nndalal@umich.edu
🖥 https://nanddalal.github.io/

Experience

Dec 2021–Present **Software Engineer, Machine Learning**, *Google*.

Oct 2019–Dec 2021 **Senior Machine Learning Engineer**, *Nines*.

Oct 2018–Oct 2019 **Machine Learning Engineer**.
Startup applying machine learning to improve radiology.
- Built FDA cleared algorithms for detecting emergent conditions on CT head scans and measuring lung nodules on CT chest scans.
- Iteratively improved computer vision (classification, detection, segmentation) and NLP (classification, summarization) models through evaluation, error analysis, radiologist feedback, and labeling.
- Designed and developed systems for using data from the Nines teleradiology clinic for analytics and model development in HIPAA compliant environment.
- Developed multimodal method of using NLP on radiology reports to study generalization properties of image based models.
- Incorporated ideas from research papers such as hybrid 2.5D detection models for false positive reduction and clustering embeddings to identify hidden stratification.
- Established patterns and processes within the engineering team for model development, deployment, and monitoring.

Feb 2017–Sept 2018 **Senior Research Engineer**, *Clarifai*.

Feb 2016–Feb 2017 **Research Engineer**.
Startup building a machine learning platform for a variety of computer vision applications.
- Improved Clarifai's general model (which predicts over 11k concepts) thus improving downstream services which use its embeddings (for transfer learning and visual search).
- Worked on transfer learning service allowing users to rapidly train image classification models on their own labeled datasets.
- Developed solutions for customers across many verticals (ear disease classification, aerial imagery detection, fashion snap and search).
- Transitioned entire company from an internal neural network toolkit to TensorFlow and maintained research toolkit used to train and serve computer vision models.
- Scaled training infrastructure from single gpu to multi-gpu/node training.
- Extended serving infrastructure to run in multiple environments (cloud, on-premise, and mobile).

Aug 2014–Dec 2015 **Research Assistant**, *University of Michigan*.
Deep Learning research advised by Professor Honglak Lee.
- Applied convolutional autoencoders to non-linear electromagnetic inverse scattering, improving reconstructions and reducing runtimes of through-the-wall imaging.
- Model improvements: experimented with adaptive learning rates and multi-task learning methods in order to scale up algorithm to harder datasets.
- System improvements: integrated matconvnet for GPU acceleration in matlab then ported code to Theano for further speedups.
- This work was presented as an abstract at IEEE AP-S/URSI 2016 in Puerto Rico.
- Worked through Professor Lee's advanced tutorials on deep learning algorithms.

Sept 2015–Dec 2015 **Research Assistant**, *University of Michigan*.

Graph Mining research advised by Professor Danai Koutra.
- Worked on parallelizing and scaling up VoG, a graph summarization algorithm developed by Professor Koutra.
- Ported over core components of VoG from original matlab codebase to python in order to leverage better libraries for multiprocessing.
- Profiled and benchmarked algorithm in order to find bottlenecks and then proposed new parallelizable subgraph generation subroutines.
- Examined the performance of these new methods by running experiments on multiple large graphs and comparing runtime, compression rate, and structure distribution.

Jan 2015–May 2015 **Corporate R&D Intern**, *Qualcomm Research*.

Deep Learning Research Team
- Developed a neural network framework for the Adreno GPU.
- Implemented core neural network internals from scratch in OpenCL.
- Profiled and benchmarked OpenCL kernels in order to incrementally improve runtime and increase utilization.
- Worked directly with the GPU architecture team to understand how to improve performance based on the hardware.
- Presented this work at an internal GPU conference and later integrated into Qualcomm's heterogeneous neural processing engine, SNPE.

May 2013–Aug 2013 **Data Science Intern**, *Bookshout*.

Bookshout is a social and group e-reading platform.
- Developed a collaborative filtering based recommendation engine for suggesting books to users.

Jun 2011–Aug 2012 **Research Assistant**, *University of Texas at Dallas*.

High school capstone project advised by Professor David Lary at the Center for Space Sciences.
- Used neural networks to cross-calibrate HNO_3 observations from multiple satellites to quantify periods of strong denitrification in the polar stratospheres.
- Presented research paper at the Texas Junior Academy of Science at Texas A&M (2nd place awards in Earth/Space Sciences).

Education

2012–2015 **B.S.E., Computer Science Engineering**, *University of Michigan*.
- GPA: 3.673, Summa Cum Laude
- Relevant coursework: Machine Learning, Graph Mining at Scale, Operating Systems, Database Management Systems, Theoretical Statistics, Matrix Algebra
- co-re-founded MSAIL (Michigan Student AI Lab): led weekly applied machine learning tutorials for undergraduate students covering classical methods for regression, classification, and clustering culminating in neural networks.
- Member of HKN (Eta Kappa Nu) and IEEE-ACM

2008–2012 **High School**, *Greenhill School*.
- GPA: 3.747
- Kilmer Award for Excellence in Science (top science student)
- Book awards for most outstanding student: Vector Calculus & Differential Equations, AP Calculus BC, AP Chemistry, AP Physics C
- Started science capstone program where students now regularly work with professors from local universities.
- Member of Math Club, Science Olympiad, and Policy Debate teams.