# Neil Movva

Computer Vision / Deep Learning, from a systems perspective

nmovva@alumni.stanford.edu

neilmovva.com

## EDUCATION

**Stanford University**                                                Palo Alto, CA
*BS and MS in Electrical Engineering, MS GPA: 3.9*              *Sep 2015 – Jun 2019*

## EXPERIENCE

**Apple**                                                            Cupertino, CA
*Deep Learning Engineer*                                                *Aug 2019 -*

Working on applied machine learning projects, primarily in computer vision.
Focus on ultra low-latency inference.

**Built Robotics**                                                San Francisco, CA
*Robotics Engineer (Intern)*                                          *Summer 2018*

Architected a multi-modal sensor system for high reliability vehicle perception.
Developed domain-specialized neural models to operate on fused sensor data.
Extensive use of model compression to deliver real-time performance on a moving vehicle.

**Nvidia**                                                          Santa Clara, CA
*Deep Learning Architect (Intern)*                                *Summer 2016, 2017*

Microarch-level performance analysis of DL training workloads (ResNets, LSTMs).
Fine-grained optimization of DL operations on the Volta GPU to maximize throughput.
BatchNorm investigation: achieved up to 20% speedup over the public cuDNN implementation.

**Intel**                                                           Santa Clara, CA
*Hard IP Group (Intern)*                                              *Summer 2015*

Exploration of NLP and ASR techniques on the Intel Edison embedded platform.

## SPECIALIZATIONS

**Languages & Tools:** Proficient with Python and C, competent in C++ and Java. Parallel code: CUDA, OpenMP, MPI. DL frameworks: PyTorch, TensorFlow. Very comfortable with Linux.

**High Performance Computing:** I operate a small server farm (48TB iSCSI SAN + many virtualized compute nodes, networked with optical 10GbE) for my DL experiments. Familiar with modern high-speed interconnects, cluster coordination, NUMA patterns, distributed filesystems.

**Relevant Coursework:**

- **EE 367**: Computational Imaging and Display
- **CS 348**: Visual Computing Systems
- **CS 231**: ConvNets for Visual Recognition
- **CS 224**: NLP with Deep Learning
- **CS 229**: Machine Learning
- **EE 282**: Advanced Computer Architecture
- **EE 376**: Information Theory
- **CS 349**: Cloud Computing
- **CS 155**: Computer Security
- **CS 144**: Computer Networking

## PUBLICATIONS

**Dynamic Field of View in a Tomographic Light Field Display.**
Movva, N. and Ginsberg, J. (equal contribution). doi:10.5594/JMI.2018.2876072.
SMPTE Best Student Paper Award. Published in *Motion Imaging Journal* (Feb 2019).