

EDUCATION

- **Carnegie Mellon University** Pittsburgh, US
MS in Electrical and Computer Engineering; GPA 3.95/4 Jan. 2017 - May. 2018(Anticipated)
- **Related Coursework**
Machine Learning (PhD) Deep Learning (PhD) Deep Reinforcement Learning (PhD)
Statistic Machine Learning (PhD) Computer Systems (MS) Compiler Design (MS)

RESEARCH INTERNSHIP

- **Carnegie Mellon University** Pittsburgh, US
Graduate Research Assistant @ ECE May. 2017 - Aug. 2017

PROJECTS

- **Optimization of Deep Learning Inference Engine on Autopilot Devices** Sep. 2017 - Dec. 2017
 - Implemented a MobileNet inference model on Nvidia Tegra TX2
 - Optimized inference performance of detecting from 10 fps to 30 fps
- **C0 Compiler** Sep. 2017 - Dec. 2017
 - Implemented a compiler translating from source language C0 to x86-64 assembly using Swift
 - Bridged C-based lexer and parser generators from Flex and Bison to Swift in compiler's front-end
- **yNN: a Lightweight Neural Network Framework** Sep. 2017 - Nov. 2017
 - Implemented purely in python without extra build phases required
 - Automatically construct computational graph for non-recursive neural networks
- **Automated Generation of Linear Algebra Kernels** May. 2017 - Aug. 2017
 - Implemented a code generator to generate level 3 GEMM kernels
 - Achieved 90% peak performance on Haswell architecture
 - Extended the kernel for algebraic path problem
- **One-shot Footwear Recognition based on Convolutional Neural Networks** Apr. 2017 - May. 2017
 - Implemented an AlexNet-like network to handle augmented images
 - Developed multiple augmentation methods to overcome limited training set
 - Achieved 99.2% accuracy on test set
- **CUDA accelerated PCA** Apr. 2017 - May. 2017
 - Achieved 200x speedup compared to CPU baseline version
- **Linux Shell** Mar. 2017
 - Supported multiple child processes and background/foreground switching
 - Eliminated race conditions between processes
- **3D reconstructions of Light Field Images** Jun. 2016 - Dec. 2016
 - Implemented a software which takes image captured by light field camera into depth adjustable images
 - Calibrated driver parameters of Lytro light field camera

SKILLS

- **Machine/Deep/Reinforcement Learning::** Tensorflow, Caffe, PyTorch, Theano, OpenAI Gym
- **High Performance/Large Scale Computing:** Intel AVX, CUDA, OpenMP, Hadoop, TensorRT
- **Programming Languages:** C++, Python, C, MATLAB, Swift, Java, OCaml, x86-64 asm