Markus Woodson

markus.woodson1@gmail.com
412-540-0490

Current graduate student looking for a place where I will learn something new everyday and be pushed by my peers to be better. Always looking for a challenge and difficult problems to solve. Research interests include deep learning, unsupervised learning, and how they can be used to help process video data

Work Experience

Summer 2016

AiCure - Computer Vision Research Intern
Applied computer vision methods to the task of pill recognition in low resolution images with obstruction of pills. Also applied deep learning and one-shot learning concepts in the practical area of pill recognition with very little images

Spring 2016

Carnegie Mellon Cylab - Research Intern
Research and development in using deep learning for full scene pedestrian detection without sliding window or cascade

Summer 2015

Hudl - Data Science Intern
Developed analysis tools for coaches and players in soccor games using neural networks, logistic regression and expectation maximization. Used tools such as Spark and Hadoop to handle big data.

Teaching Experience

Fall 2016

Mathematical Background for Machine Learning Teaching Assistant
Responsible for creating lecture material, developing and creating homeworks, and holding weekly office hours.

Fall 2014

Computational Perception Teaching Assistant
Developed homeworks, graded assignments and quizzes and held weekly office hours

Education

2016 - Current

MS Electrical and Computer Engineering - Carnegie Mellon University

2012 - 2016

BS Electrical and Computer Engineering - Carnegie Mellon University

Skills

- Python
- C++
- CUDA
- Spark / Hadoop
- MPI / Open MP
- Theano
- Torch
- scikit-learn

Projects

Attention Based Hierarchical RNN Model for Action Recognition

Implemented a hierarchical RNN model that exploits multiple temporal scales. Integrated with this was a hierarchical attention mechanism which learned to focus further as we go up in the hierarchy. I applied this model to both video captioning and action recognition tasks and got near state-of-the-art performance.

Security Quadcopter

Using gait and facial recognition, my group and I programmed a quadcopter to recognize people by their face and by the way they walk.

3-Stream Recurrent Neural Network for Action Recognition

Developed a 3 stream approach to action recognition in videos. Combined both frame-level and temporal information by using optical flow features and RNN in one of the streams. We employed cropping and sub-sampling of features in such a way to not decrease performance but save computation time.

Improving Neural Network Training Speed

Got near 130X speedup in deep neural net training using FFT and parallelization for Intel CPU's. Had half of the memory usage compared to similar FFT approaches at time of development.

Kaggle WhatsCooking Competition

Placed in the top 150 for the WhatsCooking Kaggle competition. Final model was a wide and shallow neural network using TFIDF features from stemmed ingredient lists.