

Jonathan Mitchell

Computer Vision Engineer

linkedin.com/in/jonathancmitchell
email: Jmitchell1991@gmail.com

website: jonathancmitchell.github.io
github.com/jonathancmitchell

About

I'm a current Masters CS student at UCLA with a focus in Computer Vision and Pattern Recognition. I was previously a Computer Vision engineer @Octi.tv (<https://octi.tv>), where I worked extensively on human pose estimation, instance segmentation, and gesture recognition by training neural networks to perform those tasks. My research interests are in 3D reconstruction, pose estimation, instance segmentation, cognition, causality and parallel computing. In my spare time I enjoy lifting weights, playing video games, fishing, and building computers. Please see my website for demos of my research projects.

Technical Skills

Science: C • C++ • Python • Caffe • Keras • Tensorflow • OpenCV • MATLAB • Arduino/RaspberryPi/Tiva • Solidworks • NumPy • SciPy • CUDA • Pandas • Scikit-learn • Linux • R • HFSS

Full-Stack: JavaScript (jQuery, Angular, React, Redux, Backbone, ES6, ES7) • HTML5/CSS3 • Angular • Sass • MySQL • NodeJS • MongoDB • KnexJS • Grunt • Gulp • Babel • Heroku • Git • MongoDB • TDD/Mocha/Jasmine

Deep Learning: Decision Trees • Convolutional Neural Networks • Object Detection • Semantic Segmentation • Scene Understanding • Numerical optimization for mobile integration (iOS)

Research

Instance Segmentation:

Bounding Box and Pose Detection combined

07/2018 - 09/2018

- Added a human pose detection branch to an existing Bounding Box Embedding network in order to predict human pose and instance masks from a random image. Achieve $mAP@IoU=50 = 0.559$ on COCO keypoint detection task and $AP@IoU=50 = 0.338$ on COCO Instance Segmentation benchmark.
- Re-designed data generation in order to include human keypoints, wrote entire code base: designed, trained, and tested on Microsoft's COCO dataset.
- Network is mobile-ready and can run in real time on an iPhoneX's GPU.

Bounding Box Embedding for Single Shot Person Instance Segmentation

06/2018 - 07/2018

<https://arxiv.org/abs/1807.07674>

- Wrote and designed the training pipeline, loss functions, data generation, evaluation. Trained and benchmarked on the COCO dataset. Implemented using Keras and Tensorflow.
- Implemented several different network architectures during training. Encoders include ResNet50, ResNet101, MobileNet, MobileNetV2 and decoders include U-Net, DeepLabV3+, and naive upsampling.
- Built Atrous Spatial Pyramid pooling layers in order to increase semantic information flow while maintaining receptive field view between our encoder and decoder
- Achieved 0.374 AP on the Person class compared to PersonLab's 0.382 AP, and our network is mobile ready

Semantic Instance Segmentation via Metric Learning:

03/2018 - 04/2018

- Created a neural network that branches off of a Convolutional Pose Machine to identify individual instance masks in an image
- Network is based on metric learning with a discriminative loss function that decodes semantic instances of humans and is able to identify individual human masks in a scene
- This network is different and an alternative to Mask-RCNN implementations and it's able to run on mobile (iOS)
- Designed a contrastive loss function and an encoding – decoding pipeline using RMPE server, Keras, and Tensorflow

Mobile-Mask-RCNN:

04/2018 - 06/2018

- Created a mobile-optimized version of Mask-RCNN by swapping out ResNet50 and using a MobileNet V1 or MobileNet V2 backbone network with depthwise-separable convolutions or inverted residual blocks.
- Reduced computation time to 0.07s at inference on an Nvidia Geforce Gtx 1080Ti.
- Created Feature Pyramid Network, custom ROI Alignment layers, hyper parameter tuning, mobile quantization, and weights pruning
- Project implemented using Tensorflow and Keras

Human pose:

Deep3D: Human skeletal 2D -> 3D reconstruction network for depth and missing joint inference **10/2017**

- Created a synthetic data generation pipeline to select virtual cameras for perspective projection locations and generate indefinite amounts of training data using the Panoptic pose dataset and the Human3.6M dataset
- Designed and implemented the network's preprocessing and postprocessing pipeline using Numpy, OpenCV, and Keras

2D human joint detection networks: **09/2017**

- Optimized Convolutional Pose Machines that take an image as input and extract human 2D joints
- Utilized depthwise-separable convolutional layers for network compression and mobile GPU optimization, achieving a 15x computation reduction
- (in progress) Implement custom deformable convolutional layers and deformable RoI pooling layers for high GPU utilization and Tensorflow integration

Scene recognition:

Deep Depth: **02/2018**

- Design a neural network to produce depth maps from indoor scenes in order to reject background noise when performing human pose detection
- Implement a multi-scale neural network that uses global context, gradient estimation, and refined feature smoothing for depth estimation in Keras

Breast Cancer Analysis: Malignant vs. Benign classifier | link: <http://bit.ly/2yzgICS> **05/2017**

- Extracted viable features such as smoothness, radius, symmetry, compactness and fractal dimensionality and trained machine learning models to predict whether a certain breast cancer tumor is malignant or benign
- Benchmarked and trained Naïve Bayes, SVM, and Decision Tree classifiers
- Obtained 92% accuracy on class predictions using a support vector machine

Education

Masters of Science: Computer Science **09/2018 - Present**
University of California, Los Angeles
Los Angeles, CA

Focus in Computer Vision and Pattern Recognition

Self-Driving Car Engineering Nanodegree **11/2016 - 08/2017**
Udacity
Mountain View, CA

Advanced Software Engineering Immersive **06/2016 - 10/2016**
Hack Reactor
Los Angeles, CA

Bachelor of Science: Electrical Engineering **09/2015 - 12/2015**
University of California, Davis
Davis, CA

Focus in signal processing and microwave engineering

Awards

University of California Davis: Dean's Honors List: July 2015, December 2015

Publications

Bounding Box Embedding for Single Shot Person Instance Segmentation | link: <https://arxiv.org/abs/1807.07674>

Autonomous Vehicle Speed Estimator from Dashboard Cam | link: <http://bit.ly/2n9gPwQ>

Behavioral Cloning | link: <http://bit.ly/2mblzxl>

Work Experience

Computer Vision Engineer **06/2017 - 09/2018**
SuperMediaFuture dba Octi Inc. | link: <http://www.octi.tv> **Los Angeles, CA**

- Implement, design, and experiment with deep neural networks using Caffe, OpenCV, Tensorflow, and Keras
- Focus on Instance Segmentation, Human Pose estimation and 3D reconstruction. (research section above)
- Research, design and optimize 3D reconstruction networks for mobile (iOS)
- Design, train, and test metric-based instance segmentation networks to extract individual segmentation masks for each human in an image. This is a non-bounding box approach that utilizes pixel embeddings and an encoding – decoding network with a discriminative loss function (different from mask rcnn). This network is also optimized for mobile integration (iOS).

Research Volunteer **03/2017 - 05/2017**
University of Southern California **Los Angeles, CA**

- Focus on unsupervised and semi-supervised learning, object segmentation and convolutional network design

Technical Mentor **12/2016 - 06/2017**
Udacity **Los Angeles, CA**

- Technical mentor providing support to 17 students in Udacity's Self-Driving Car Nanodegree program
- Science Tutor** **11/2016 - 05/2017**
Hayutin Associates **Los Angeles, CA**
- Math, physics, and programming tutor for high school and college students
- Founder** **01/2016 - 07/2016**
Pushlens **Los Angeles, CA**
- Invented a contact lens insertion device so individuals can insert contact lenses without having to physically touch them
 - Designed product in Solidworks, performed extensive flow simulation, stress testing and 3D printed prototypes
- Automation Engineering Intern** **06/2015 - 09/2015**
Genentech **Oceanside, CA**
- Increased batch checking precision and efficiency through improving gas chromatography review tools
- Software Development Engineering Intern** **06/2014 - 09/2014**
inSightec
- Created algorithms that contributed towards Parkinson's tremor's cure through designing a feedback control system that monitors an MRI-guided focused-ultrasound wave during a neurosurgical procedure
 - Designed control system for real-time simulation of musculoskeletal responses in MATLAB

Open Source

- Contributor**
- Keras team** | link: <https://keras.io/> **05/2017**
- Official implementer of MobileNetV2 (<https://arxiv.org/abs/1801.04381>) for Keras. MobileNetV2 is a very light weight general backbone network that can be used for various tasks such as object localization, classification, and segmentation. It has been pre-trained on imagenet.
 - MobileNetV2 dramatically reduces both parameter count and computational cost. Specifically, for the task of object detection on ImageNet dataset, MobileNetV2 has 0.8M fewer params and outperforms MobileNetV1 by 0.6 %.
 - Network was merged on 05/07/2018 (#10047) and can be found under keras_applications: https://github.com/keras-team/keras-applications/blob/master/keras_applications/mobilenet_v2.py

Personal Projects

- Vehicle and Lane Line Detection** | vehicle-repo: <http://bit.ly/2qTMIrG> lane-repo: <http://bit.ly/2oK2hjj> **03/2017**
Detects vehicles and lane lines from a video stream using a support vector machine and computer vision techniques
- Designed cell-based window search method to extract image features and predict vehicle locations using OpenCV
 - Implemented adaptive color thresholds, filtering, windowing, and morphological operations to detect lane lines
- Speed Estimation – ML Eng** | <http://bit.ly/2n9gPwQ> | github.com/jonathancmitchell/speedchallenge **03/2017**
Neural network dashboard cam speed estimator for autonomous vehicles trained on 8616 images extracted from a video
- Utilized dense optical flow analysis and architected conv-net processing pipeline to achieve MSE of ~5.6
- Behavioral Cloning – Self-Driving Car Engineer** | <http://bit.ly/2mblzxl> **02/2017**
Trained a model to simulate a self-driving car by predicting steering angles while driving in a unity simulation
- Normalized a multi-modal data distribution by performing data augmentation and creation
 - Created a creative preprocessing pipeline with brightness augmentation, vertical flipping, resizing and cropping
- Traffic Sign Classification – Machine Learning Engineer** | <http://bit.ly/2kkFlzN> **01/2017**
Classify Traffic Signs from more than 40 classes in a German Traffic Sign Dataset of 50,000 32x32 images
- Implemented a LeNet-like CNN using Google's Tensorflow deep learning framework
 - Utilized normalization, data creation, regularization methods to increase accuracy to ~ 94.5% at test using Tensorflow
- GamesReviewAPI – Full Stack Engineer** | <https://gamereviewdocs.herokuapp.com/> **11/2016**
A restful API that allows users to search through 18k game reviews and query for top scores by genre, platform, and title

- Designed ES7 generator functions to synchronously seed the database with 18k row entries

Radar Sensor – Electrical Engineer | <http://bit.ly/2fNKu9w>

02/2014 - 06/2014

24 Ghz Radar sensor to detect real time speed and range measurement

- Programmed a micro-processor to control a patch antenna array and transmit signals to antenna transceiver

Conferences Attended

Computer Vision and Pattern Recognition	06/2018
Computer Vision Foundation	Salt Lake City, Utah
AI Vision link: http://ai.vision/	03/2017
AI By The Bay	San Francisco, CA
Self-Driving Cars link: http://self.driving.cars/	03/2017
AI By The Bay	San Francisco, CA
TechCrunch Disrupt	09/2015
TechCrunch	San Francisco, CA
TechCrunch Disrupt	09/2014
TechCrunch	San Francisco, CA

Leadership

Founder and Lecturer	01/2017 - Present
Self-Driving Cars Los Angeles meetup.com/Los-Angeles-Self-Driving-Car-Meetup/	Los Angeles, CA
• Prepare and present lectures on machine learning, computer vision, and image processing to	385 members
Student Body Senator	12/2014 - 12/2015
Associated Students Of University of California Davis	Davis, CA
• Oversaw 26 student-run businesses on campus and held responsibility over all revenues and property	
• Co-managed an annual 11.8M budget and merged two student-run businesses to save another	
• Administrated hiring efforts and conducted interviews for over 20 student-run business managers	
President	09/2014 - 03/2015
Society of Biological Engineers at UC Davis	Davis, CA
• Managed club related finances, facilitated meetings and increased membership from 28 to 53 members	

Research interests

Instance Segmentation:

- Top down and bottom up approaches to human instance segmentation

Human Pose Estimation:

- 2D keypoint localization
- 3D reconstruction
- Structure from motion
- Gesture recognition

Deep learning:

- Convolutional Pose Machines
- Reinforcement learning
- Deep residual based networks
- Faster RCNN, SSD, Mask RCNN, Instance Segmentation
- Capsule networks

Medical Imaging:

- MRI-guided focused ultrasound therapy

Autonomous Vehicles:

- Path planning and localization

Multi-view Geometry:

- Synthetic data generation using virtual camera selection
- 3D scene reconstruction

Mathematics:

- Algebraic Geometry, Julia and Mandelbrot sets, bifurcation and chaos theory

Haskell, CUDA programming and operating systems

Relevant Coursework

UCLA

Los Angeles, CA, 09/2018 - Present

- CS 266: Statistical modeling for vision and cognition
- CS 276A: Computer Vision and Pattern Recognition
- CS 174A: Computer Graphics
- CS231n – Convolutional Neural Networks for Visual Recognition

Stanford | online, 11/2016 - 02/2017

Udacity | online, 11/2016 - 08/2017

Self-Driving Car Nanodegree term 1

- Computer Vision
- Deep Learning

Los Angeles, CA, 07/2016 - 11/2016

Hack Reactor

- Database schema design and system architecture with MySQL,
- Backend development with NodeJS and ExpressJS
- Frontend development with ReactJS, AngularJS, BackboneJS, HTML, CSS, JavaScript
- Data structures and algorithms

UC Davis:

Davis, CA, 2009 - 2015

- Probability and statistics
- Fluid mechanics
- Electronic circuits and systems
- Control systems
- Semiconductor device fundamentals I and II
- Digital signal processing
- Calculus I, II, II, vector calculus, linear algebra and differential equations
- Classical and modern physics
- Bioinstrumentation and control
- Antenna design and radio / microwave engineering
- Electromagnetism
- C programming and problem solving
- Semiconductor fabrication

What I am reading at night:

Multi-View Geometry for Computer Vision – Hartley and Zisserman

Programming Massively Parallel Processors – Kirk and Hwu

Effective C++ – Meyers

3D Math Primer for Graphics and Game Development – Dunn and Parberry

References

References available upon request.
