

# Jonathan Mitchell

linkedin.com/in/jonathancmitchell  
email: jcmitchell@g.ucla.edu

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

---

## About

I'm a current PhD CS student at UCLA with a focus in Computer Vision and Pattern Recognition. My interests lie at the intersection of Computer Vision and Security - providing security defenses in safety critical AI systems including Facial Recognition and Autonomous Driving. I was previously a Computer Vision engineer @Octi.tv where I worked extensively on human pose estimation, instance segmentation, and gesture recognition by training neural networks to perform those tasks.

---

## 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

**Stochastic Security: Adversarial Defense Using Long-Run Dynamics of Energy-based models** 07/04/2020

- <https://arxiv.org/abs/2005.13525> submitted to ICLR2021.
- Constructed State of the Art Adversarial Defense preprocessing algorithm without adversarial training 55.5 acc

**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 in the wild. Achieved  $mAP@IoU=50 = 0.559$  on Microsoft MSCOCO 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 and 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 a15x 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/2yzglCS>

**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

### Doctor of Philosophy: Computer Science

**09/2018 - Present**

University of California, Los Angeles

**Los Angeles, CA**

Focus in Computer Vision and AI Security

### Masters of Science: Computer Science

**09/2018 - 06/22/2020**

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

---

## Work Experience

### Intern - Autonomous Driving

**06/2020 - 09/2020**

Nvidia

**Santa Clara, CA**

- Shadow mode active learning loop for Nvidia drive platform
- NCAP validation for perception based systems for verification and robustness

### Graduate Student Researcher

**07/2019 - Present**

### UCLA Center for Vision Language Cognition and Autonomy

**Los Angeles, CA**

- Member of UCLA research group working on Computer Vision, Energy Modeling, MCMC, AI Security

### 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**

<b>University of Southern California</b>	<b>Los Angeles, CA</b>
<ul style="list-style-type: none"> <li>Focus on unsupervised and semi-supervised learning, object segmentation and convolutional network design</li> </ul>	
<b>Technical Mentor</b>	<b>12/2016 - 06/2017</b>
<b>Udacity</b>	<b>Los Angeles, CA</b>
<ul style="list-style-type: none"> <li>Technical mentor providing support to 17 students in Udacity's Self-Driving Car Nanodegree program</li> </ul>	
<b>Science Tutor</b>	<b>11/2016 - 05/2017</b>
<b>Hayutin Associates</b>	<b>Los Angeles, CA</b>
<ul style="list-style-type: none"> <li>Math, physics, and programming tutor for high school and college students</li> </ul>	
<b>Founder</b>	<b>01/2016 - 07/2016</b>
<b>Pushlens</b>	<b>Los Angeles, CA</b>
<ul style="list-style-type: none"> <li>Invented a contact lens insertion device so individuals can insert contact lenses without having to physically touch them</li> <li>Designed product in Solidworks, performed extensive flow simulation, stress testing and 3D printed prototypes</li> </ul>	
<b>Automation Engineering Intern</b>	<b>06/2015 - 09/2015</b>
<b>Genentech</b>	<b>Oceanside, CA</b>
<ul style="list-style-type: none"> <li>Increased batch checking precision and efficiency through improving gas chromatography review tools</li> </ul>	
<b>Software Development Engineering Intern</b>	<b>06/2014 - 09/2014</b>
<b>inSightec</b>	
<ul style="list-style-type: none"> <li>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</li> <li>Designed control system for real-time simulation of musculoskeletal responses in MATLAB</li> </ul>	
<b>Open Source</b>	
<b>Contributor</b>	
<b>Keras team</b>   link: <a href="https://keras.io/">https://keras.io/</a>	<b>05/2017</b>
<ul style="list-style-type: none"> <li>Official implementer of MobileNetV2 (<a href="https://arxiv.org/abs/1801.04381">https://arxiv.org/abs/1801.04381</a>) 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.</li> <li>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 %.</li> <li>Network was merged on 05/07/2018 (#10047) and can be found under keras_applications: <a href="https://github.com/keras-team/keras-applications/blob/master/keras_applications/mobilenet_v2.py">https://github.com/keras-team/keras-applications/blob/master/keras_applications/mobilenet_v2.py</a></li> </ul>	
<b>More Projects</b>	
<b>Vehicle and Lane Line Detection</b>   vehicle-repo: <a href="http://bit.ly/2qTMIrG">http://bit.ly/2qTMIrG</a> lane-repo: <a href="http://bit.ly/2oK2hjJ">http://bit.ly/2oK2hjJ</a>	<b>03/2017</b>
<i>Detects vehicles and lane lines from a video stream using a support vector machine and computer vision techniques</i>	
<ul style="list-style-type: none"> <li>Designed cell-based window search method to extract image features and predict vehicle locations using OpenCV</li> <li>Implemented adaptive color thresholds, filtering, windowing, and morphological operations to detect lane lines</li> </ul>	
<b>Speed Estimation – ML Eng</b>   <a href="http://bit.ly/2n9gPwQ">http://bit.ly/2n9gPwQ</a>   <a href="https://github.com/jonathancmitchell/speedchallenge">github.com/jonathancmitchell/speedchallenge</a>	<b>03/2017</b>
<i>Neural network dashboard cam speed estimator for autonomous vehicles trained on 8616 images extracted from a video</i>	
<ul style="list-style-type: none"> <li>Utilized dense optical flow analysis and architected conv-net processing pipeline to achieve MSE of ~5.6</li> </ul>	
<b>Behavioral Cloning – Self-Driving Car Engineer</b>   <a href="http://bit.ly/2mblzxl">http://bit.ly/2mblzxl</a>	<b>02/2017</b>
<i>Trained a model to simulate a self-driving car by predicting steering angles while driving in a unity simulation</i>	
<ul style="list-style-type: none"> <li>Normalized a multi-modal data distribution by performing data augmentation and creation</li> <li>Created a creative preprocessing pipeline with brightness augmentation, vertical flipping, resizing and cropping</li> </ul>	
<b>Traffic Sign Classification – Machine Learning Engineer</b>   <a href="http://bit.ly/2kkFlzN">http://bit.ly/2kkFlzN</a>	<b>01/2017</b>
<i>Classify Traffic Signs from more than 40 classes in a German Traffic Sign Dataset of 50,000 32x32 images</i>	
<ul style="list-style-type: none"> <li>Implemented a LeNet-like CNN using Google's Tensorflow deep learning framework</li> <li>Utilized normalization, data creation, regularization methods to increase accuracy to ~ 94.5% at test using Tensorflow</li> </ul>	

**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

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

---

#### Leadership

<b>Founder and Lecturer</b>	<b>01/2017 - Present</b>
<b>Self-Driving Cars Los Angeles   <a href="https://meetup.com/Los-Angeles-Self-Driving-Car-Meetup/">meetup.com/Los-Angeles-Self-Driving-Car-Meetup/</a></b>	<b>Los Angeles, CA</b>
<ul style="list-style-type: none"> <li>• Prepare and present lectures on machine learning, computer vision, and image processing to 385 members</li> </ul>	
<b>Student Body Senator</b>	<b>12/2014 - 12/2015</b>
<b>Associated Students Of University of California Davis</b>	<b>Davis, CA</b>
<ul style="list-style-type: none"> <li>• Oversaw 26 student-run businesses on campus and held responsibility over all revenues and property</li> <li>• Co-managed an annual 11.8M budget and merged two student-run businesses to save another</li> <li>• Administrated hiring efforts and conducted interviews for over 20 student-run business managers</li> </ul>	
<b>President</b>	<b>09/2014 - 03/2015</b>
<b>Society of Biological Engineers at UC Davis</b>	<b>Davis, CA</b>
<ul style="list-style-type: none"> <li>• Managed club related finances, facilitated meetings and increased membership from 28 to 53 members</li> </ul>	

---