

Phillip Kuznetsov

2311 Channing Way Apt.C
Berkeley, CA 94704
801-971-4414
philkuz@berkeley.edu
github.com/philkuz

EDUCATION

University of California Berkeley - Undergrad, EECS 2018

Relevant Courses: Computational Photography, Graphics, Computer Vision, Machine Learning, Probability and Stochastic Processes, Algorithms, Convex Optimization

EXPERIENCE

Berkeley AI Research

ADVISOR: ALYOSHA EFROS; FEBRUARY 2018- PRESENT

Currently working on a Deep Learning system to determine whether a Music-Video is good or not. Dataset of human ratings for music-video pairs collected via Mechanical Turk. Architectures experimented with so far include I3D for video representation and VGGish for audio representation.

Alinea AI

MARCH 2018 - PRESENT

Deep learning and data science consultancy. Previous projects include predicting movements in cryptocurrency markets and predicting customer behavior using financial securities data.

Machine Learning at Berkeley - President

DECEMBER 2015 - MAY 2018

Founded Berkeley's first machine learning club. Grew the organization to 80 members ranging from Freshman to PhD students. Through the club, I've also [conducted research](#), [taught courses](#) and lead [workshops](#) on topics like "[Making Art with Deep Learning](#)". All project/event repos located at github.com/mlberkeley.

Adobe Research

ADVISOR: BRYAN RUSSELL; MAY 2017 - AUGUST 2017

Extended a video action recognition deep learning model for Adobe Stock Videos to create a better feature representation for Adobe Stock's search functionality.

Amazon Alexa Prize

ADVISOR: JOHN DENERO; NOVEMBER 2016 - AUGUST 2017

One of the sponsored teams entered in the inaugural Alexa Prize competition. Given a \$100k research grant and unlimited to AWS to build a conversational chatbot. See [proceedings paper](#).

PAST RESEARCH PROJECTS

CANs - ML@B

Implementation and extension of Creative Adversarial Networks - [Repo](#) received 62 stars on Github. Future work planned to extend the model to use the WGAN-GP objective function and then Progressive Growing of GANs.

Transferability of Adversarial Attacks in the MAML Framework

We propose a series of experiments designed to test the susceptibility of [MAML](#) to adversarial attacks. Paper. Accepted to the 2017 [Deep Learning Security Workshop](#) in Singapore chaired by Prof. Dawn Song.

LANGUAGES

Python
Java
JavaScript
HTML / CSS
C++
C#

TECHNOLOGIES

Python
Java
JavaScript
HTML / CSS
TensorFlow
Caffe
Docker
AWS/Google Cloud
Mechanical Turk

PUBLICATIONS

Artificial Intelligence Safety and Security Textbook

Chapter on Adversarial Machine Learning

First author on Adversarial Machine Learning chapter. To be published in August 2018. [Preprint chapter](#).

O'Reilly Blog - 2017

Generating Images with PixelRNNs.

O'Reilly article on using TensorFlow to generate novel images using PixelRNNs. Written to be easily accessible for newcomers to TensorFlow. [Link to article](#).