

ADRIAN YI

a@adrianyi.com | 267-210-4976 | San Francisco, CA 94107

[Personal page](#) | [LinkedIn](#) | [GitHub](#)

ML Engineer with experience in programming and deep learning

SKILLS

Programming	Python, HTML/CSS/JavaScript, Unix, MATLAB
ML	Classification, Regression, NLP, Computer Vision, Sklearn, TensorFlow, Keras, PyTorch
Analysis	SQL, Pandas, Scipy, Matplotlib, Seaborn, D3.js, Statistics, Data wrangling
Deployment	AWS, Flask, Docker, PostgreSQL, Heroku

EXPERIENCE

ML Engineer – Clusterone Inc. June 2018 – Present

- Improved video quality loss prediction using deep learning and decreased MSE by an order of magnitude for a video compression company
- Implemented distributed XGBoost and Horovod on Clusterone platform using OpenMPI
- Authored a blog on distributed TensorFlow training that was one of the top Google search results
- Created and maintained Docker images used for Clusterone platform

Fellow – The Data Incubator March 2018 – May 2018

- Completed weekly projects with real life data that covered topics like data wrangling, machine learning, natural language processing, and distributed processing (i.e. MapReduce/Spark)
- Built an NLP project and was selected to present to industry data scientists

Research Fellow & Consultant – Boston University May 2011 – February 2018

- Led a team of 3 researchers to create models that explains the inner workings of protein used to control neuron cell activity
- Automated data pipeline and analysis using MATLAB, reducing analysis times by over 10x
- Built the lab's first time-resolved analysis technique using a custom baseline correction and SVD that maps the protein's state from 10^{-9} to 10^3 seconds range
- Published 5 papers and presented posters at 2 conferences

PROJECTS

Topic Insight

- Designed a general NLP product that finds similar documents and extracts words correlated with business metrics like higher views, faster responses for emails, etc.
- Deployed a live app using 260,000 grant abstracts from 5 data sources, each with millions of rows
- Tech: Python, NLP, LDA, Pandas, Scikit-learn, Flask, Docker, PostgreSQL, visualization

Dog Breed Classifier App

- Built an app that can classify an image into one of 133 dog breeds with 86% accuracy
- Deployed the model on AWS using Flask and Docker
- Tech: Python, CNN, Keras, transfer learning, AWS, JavaScript

EDUCATION

Boston University Boston, MA | September 2010 – May 2017

- Ph.D. and M.A., Physics
- Teaching Fellow - Intro to Physics for engineers (2 years)

Franklin and Marshall College Lancaster, PA | September 2005 – May 2009

- B.A. Physics and Mathematics