

Joseph Chuang

jochuang@stanford.edu (607)-261-2174 jochuang.com github.com/josephch405 [in/jochuang](https://in.jochuang)

Education

Stanford University, MS in Computer Science – Palo Alto, CA

Fall 2019 – Spring 2021

Cornell University, BS in Computer Science – Ithaca, NY

Fall 2016 – Spring 2019

GPA: 4.0; Relevant coursework: Algorithms, Machine Learning and Intelligent Systems, Functional Programming, Advanced Topics in ML, Advanced ML Systems, Operating Systems, Stochastic Processes

TA Experience: Machine Learning, Algorithms

Experience

Software Engineering Intern - Google, Sunnyvale, App Maker

May – Aug 2019

- Designed and wrote data model field recommendation system using MapReduce and SVD
- Implemented Magic Magnets feature, segmenting photos of draft webpage layouts with Google Cloud Vision AutoML and generating hierarchical widget layouts with beam search

Research Intern - Taiwan AILabs, Taipei, Taiwan

Dec 2018 – Jan 2019

- Adapted VoiceFilter technique for denoising audio spectrograms with Pytorch CNN/LSTMs
- Dockerized and deployed service for K8 deployment and integration with Chinese ASR pipeline

Software Engineering Intern - Asana, New York City

Sep – Nov 2018

- Drove data model migration for core product; redesigned Typescript client, then built distributed migration in Node.js, enabling project type conversion between Boards and Lists views
- Worked on ETL migration from Redshift to Spark by working with Airflow and Scala processors

Software Engineering Intern - Google, Sunnyvale, Cloud Shell Tutorials

May – Aug 2018

- Designed and implemented tutorial recommendation service by processing corpus via TF-IDF and NMF, then serving based on author-provided tags

Distributed Facial Recognition Lead - Cornell Data Science Project Team (CDS)

Feb – May 2019

- Set up and configured on-prem bare metal Kubernetes cluster of four nodes
- Oversaw implementation of test Facial Recognition project consisting of frontend Node.js pod and distributed backend Python inference pods, leveraging features such as volume provisioning and horizontal autoscaling

Kaggle Subteam Lead - Cornell Data Science Project Team (CDS)

Aug – Dec 2017

Projects

Paranet - Academic Research Project

Oct – Dec 2017

- Designed and evaluated novel deep architecture, an extension from DenseNet with novel cascaded skip connections allowing early inference at varying stages for practicality
- Ran Pytorch experiments involving logit matching and alternate parameter sharing architectures

Machine Learning Playground - ml-playground.com

Jun – Aug 2017

- Designed a visual playground featured on the front page of Product Hunt, highlighting differences between ML algorithms through end-device data generation, training and inference
- Implemented library of ML algorithms (KNN, decision trees, neural nets etc) with JS, math.js

Skills

- Data Science and Machine Learning: Python, Numpy, Pandas, Scikit-Learn, Pytorch, Tensorflow
- Devops: Docker, Kubernetes, Google Cloud Platform
- Web Development: JavaScript, TypeScript, Node.js, React, Redux, Firebase, Express.js
- Software Development: Java, C, C++, Golang, Protobuf