

Joseph Chuang

jcc436@cornell.edu (607)-261-2174 jochuang.com github.com/josephch405 [in/jochuang](https://in.jochuang)

Education

Cornell University - Ithaca

Fall 2016 – Current

- Bachelors of Science in Computer Science, College of Engineering - GPA: 3.99
- Relevant coursework: Algorithms, Machine Learning and Intelligent Systems, Computer Systems, Functional Programming, Advanced Topics in ML, Advanced ML Systems, Operating Systems, Stochastic Processes

Experience

Software Engineering Intern (Current) – Asana, New York City

Sep – Dec 2018

Software Engineering Intern – Google, Sunnyvale

May – Aug 2018

- Implemented metrics pipeline for external Google Cloud Shell Tutorial authors to receive behavioral statistics in Google Analytics and Data Studio, using JS and TypeScript
- Developed a WYSIWYG authoring workflow within Google Docs by constructing a Markdown importing tool with C++, a previewing tool for walkthroughs on Golang App Engine, and a Docs plugin in Apps Script hosting the walkthrough runtime; increased the count of available official Cloud Walkthroughs by 83% with this new workflow
- Designed and implemented tutorial recommendation service on internal Java serverless platform, first by processing corpus and keywords via TF-IDF and NMF, then serving based on author-provided tags

Machine Learning (CS 4780) Teaching Assistant

Jan – May 2018

Algorithms (CS 4820) Teaching Assistant

Aug – Dec 2017

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

Aug – Dec 2017

- Led subteam of 9 in Kaggle competitions, exploring text/image datasets and developed competitive models in XGBoost, Pytorch and Tensorflow/Keras
- Directed team effort to compete in the Zillow Dataset challenge, using Pandas, Numpy, XGBoost, LGBM, Catboost and general ensembling methods to generate a top ~5% submission

Researcher - Team Ursa

Jul – Sep 2017

- Researched and engineered a solution for audio alignment tasks, using FFMPEG, Numpy and Scikit to downsample, preprocess, and predict offsets between separate audio recordings of the same source event
- Built a modular system that achieved human-level performance; deployed code as AWS Lambda serverless function

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
- Implemented and ran experiments in Pytorch involving logit matching and alternate parameter sharing architectures

Engag-Ed - Big Red Hacks Cornell - Microsoft Prize, Best UI

Sep 2017

- Directed 4-person team in building an intelligent classroom suite featuring automatic attendance tracking, classroom engagement analysis, and instant polling; implemented via RESTful Express server, Microsoft Cognitive Services Emotion/Face APIs, and backed by persistent MongoDB backend
- Won the Microsoft Grand Prize and Best User Interface out of 60 teams

Machine Learning Playground - ml-playground.com

Jun – Aug 2017

- Designed a visual playground using React and HTML5 featured on front page of Product Hunt
- Highlighted differences between ML algorithms through end-device data generation, training and inference
- Implemented library of trainable ML algorithms (KNN, decision trees, neural networks etc) with JS, math.js

SkinnerDB - Academic Research Project

Jan – May 2018

- Extended reinforcement-learning based database system by parallelizing query execution and join order exploration across 30 processors with Java stream-based lambda execution
- Provided up to 4x speed increases on select TPC-H benchmarks against original serial executor

Skills

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