

Gary Zheng

www.dongyuzheng.com • garydzheng@gmail.com • www.github.com/dongyuzheng

SUMMARY

Languages Python, SQL, C++, Java
Technologies Pandas, Jupyter, Spark, Django, Docker, Chef, Git

EXPERIENCE

Facebook, *Location Infra*, Applied Machine Learning Intern (New York, NY) *Sep 2018 – Dec 2018*

- Researched and developed proprietary semi-supervised methods and tooling for the curation of labeled training data from billions of data points
- Authored visualization software for labeled data, improving efficiency in data quality spot-checking
- Developed feature space visualizations through dimensionality reduction and clustering in order to identify data distribution mismatch

Yelp, *Search Logging & Metrics*, Data Science Intern (San Francisco, CA) *Jan 2018 – Apr 2018*

- Researched, designed, and implemented models for real-time outlier detection on streaming data
- Implemented an automated metric analysis application to report anomalies in system behaviour
- Developed Spark drivers to perform grid search and random search for parameter sweeping
- Authored a JupyterLab extension for submitting Jupyter Notebooks to Review Board for peer review

Yelp, *Reservations*, Software Engineering Intern (San Francisco, CA) *May 2017 – Aug 2017*

- Implemented full-stack A/B tests and performed analysis on their effects to user metrics and SEO
- Overhauled Yelp Reservation's A/B testing framework, vastly improving experiment config process
- Managed the continuous deployment of Yelp Reservation's Elasticsearch cluster and client upgrade

NVIDIA, *Infrastructure*, Software Engineering Intern (Santa Clara, CA) *Aug 2016 – Dec 2016*

- Forked development of Packer-XenServer plugin contributing numerous features and bugfixes; currently used to automate builds for all of NVIDIA's DGX and GeForce Now server operating system images
- Authored a distributed web platform for provisioning scalable Jenkins clusters with Docker Swarm; implemented a React.js dashboard and an API backend with Express.js, Elasticsearch, and LDAP

BlackBerry, *Chef*, Infrastructure Software Developer (Waterloo, ON) *Jan 2016 – Apr 2016*

- Authored a Chef gem for managing Maven artifacts in Sonatype Nexus repositories with 1500+ users
- Pushed major features to OpenNebula Chef provisioning driver; open-sourced with 5000+ downloads
- Authored a cookbook to deploy highly-available load balancers with Keepalived and Nginx/HAProxy

PROJECTS

RezQ.io, Project Lead; Backend & Infrastructure

- Ongoing senior year capstone project: a web platform for receiving and giving resume critiques

See more projects with demos and source code at <https://dongyuzheng.com/projects/>

EDUCATION

University of Waterloo, Bachelor of Software Engineering *2014 – 2019 (expected)*

- Dean's Honours List (all terms); Cumulative average: 90.12% (3.92 GPA)
- Teaching assistant: Calculus 2 for Honours Mathematics (Winter 2017)
- Perfect score on the 2014 Canadian Computing Competition, Junior