

Relevant Work Experience

Lyft – Senior Engineering Manager, Mapping (Feb 2019 – Present)

Senior Engineering Manager, Mapping Experience (“MapXP”) & Mapping Data Engineering (Jan 2020 – present)

- Lead the MapXP team, the mobile engineering and user-facing pillar of the Mapping org consisting of 38 engineers and engineering managers based out of SF, Seattle, Kyiv, and Minsk; the team is responsible for the maps embedded in our apps; new map-based features, products, and capabilities; and the SDK that other teams can leverage
- MapXP owns LyftNav and LyftMap, which are replacing Google navigation and Google maps respectively to control the end-to-end user experience and reduce costs by ~\$60MM/yr
- Took ownership of MapXP pillar in Jan '21 and led significant changes to the team's operating mode and process to bridge an 11 hour time difference, the result of which was an exceptionally productive and successful year in which LyftNav scaled from 0 to 3 million rides and from 0 to 21 of Lyft's top 30 regions.
- Updated 2021 plan; organized, led, and tracked cross-functional 2022 planning; co-authored 3-year strategy

Engineering Manager, Localization (Feb 2019 – Dec 2020)

- Responsible for Lyft's core, high-traffic distributed systems for ingestion, storage, and serving of locations (60k+ RPS on a Python, Go, DDB, Redis, Kafka stack), passenger/rider localization, and client GPS and sensor data collection
- Managed 14 server and mobile SWEs across 3 teams (Platform, Locations Accuracy, Client), including staff engineers
- Owned '20 objective to adopt Kubernetes, improve SLA adherence, and reduce costs by 45% across the 74-engineer org
- Worked closely with research scientists to develop and execute on rider localization strategy, focusing on airports
- After 3 incidents in quick succession brought down Lyft in my first 2 months, I aggressively prioritized reliability work, revised code standards, review guidelines and our PR template, started monthly operational reviews, wrote a framework to proactively assess system risk, kicked off a reliability 'offsite,' and personally reviewed all PRs when we lacked a tech lead; as a result, we had no critical incidents in H2 of 2019 and improved availability of 83% of endpoints
- Introduced more planning discipline (e.g. grooming and a scrum master, reviewing board at standup) to address consistently missed sprint commitments and added retrospectives for continuous feedback and learning, increasing morale; and individual development plans for all reports to get and stay aligned on expectations, goals, and progress

Democratic National Committee – Staff Software Engineer (Aug – Nov 2020)

- Identified and fixed system monitoring and reliability gaps in data infrastructure, introducing industry best practices
- Documented, socialized, and led implementation of new org-wide standards for alerts, dashboards, metrics, logging, feature flags, and measuring reliability to reduce risk and make problems much easier to diagnose and faster to fix
- Added system, business, and data integrity metrics and dashboards (in Terraform for Datadog) for key pipelines (in Airflow) processing NGP VAN data, voter files, and absentee/early vote data; added a DAG to collect Airflow, BigQuery, and MySQL metrics (e.g. time from enqueueing to start, connection saturation, P50/P90 query duration)
- Built generic Python metric library leveraging StatsD and Airflow call-backs and introduced feature flag support

605 – Group VP of Technology & Engineering (Nov 2016 – Feb 2019)

Founding employee of new \$70MM company combining the engineering & analytics expertise from AMG, exclusive data rights from Spectrum, and deep industry connections to disrupt video ad market

- Oversaw the engineering, product, and IT groups, set the technical direction for the company
- Led architectural discussions, drove creation of standards, and hosted design reviews with the Director of Engineering
- Helped data team create Spark-based ETL jobs (in Scala) to ingest ~600MM rows daily of TV viewership & meta data
- Led move from Redshift to S3-based data lake (Parquet) to abstract storage from computation and reduce costs
- Oversaw the UX research and UI and feature refresh of our media planning & analysis products
- Created POCs & did hands-on development for Spark ETL & optimization engine and RoR web apps, time permitting
- Grew the team to 23 staff and ~10 contractors and managed an \$8.4MM team and infrastructure budget

The White House – U.S. Digital Service (USDS) – Digital Service Expert (Nov 2015 – Jan 2017)

Group of top tech talent recruited to use private-sector best practices to improve delivery of key government digital services

- Served variously as a technical expert, engineering lead, & management consultant depending on project needs within the Executive Office of the President (EOP) and across agencies
- Focused on homeland security projects such as the U.S. import system, the traveller experience, the refugees and immigration process, and government-wide efforts like modern collaboration tools
- Established team at Customs to help 200+ engineers improve the stability, test coverage & monitoring of the 'new' federated US import system (historically paperwork to 47 different agencies), a \$1B+ project years behind schedule; built a public availability dashboard (in Python/Flask on AWS) using AppDynamics data to improve accountability
- Improved the security of the immigration process by building a web app (in Node.js) and ETL process to rationalize the data from multiple criminal databases across departments

Analytics Media Group – VP of Engineering, Technical Co-founder (Mar 2013 – Oct 2016)

Co-founded media analytics and software start-up to bring digital-like insights to TV after identifying a gap at the DNC; grew the company to a 30+ staff and sold for more than 12x the initial investment to become 605

- Created patent-pending AI & machine-learning based optimization engine in Hadoop (MapReduce & later Spark on AWS EMR) as part of media planning & analysis software with a Rails frontend and a Rack-based micro-services layer
- Designed data model for 70+ billion rows of cable set top, consumer, voter, and related data in Redshift
- Optimized \$110MM of ad spending. Worked with Fortune 500 companies in tech (Uber), entertainment (e.g. NBC, HBO, New Regency, MSG), hospitality (e.g. Choice Hotels), and candidates for Senate & Governor
- Managed team of 10 engineers, designers, and a project manager, a personnel & data budget of \$2.0MM, and infrastructure (Amazon Web Services & local) for the full staff

Presidential Inaugural Committee – Senior Software Engineer (Dec 2012 – Feb 2013)

- Developed responsive, low-latency web app to manage volunteer signup & admin (Rails, Redis & PostgreSQL on AWS) and implemented modified stable matching algorithm to match 60k volunteers and events

Democratic National Committee – Lead Developer – Deputy Director, Media Tracking (Aug 2011 – Nov 2012)

- Led design and development of new reporting and analytics platform for media spending, polling, field and other critical campaign data used by senior staff of presidential campaign and national party committees
- Planned product with key stakeholders, designed schema, wrote complex analytics queries, and developed front & backend (Rails, PostgreSQL, Vertica, Redis, & D3.js) & led staff of four engineers

Social Security Administration – Computer Scientist and Hardy-Apfel Fellow (Jul 2010 – Aug 2011)

- Legacy mainframe-based systems prevented deployment of web-based services expected by the public
- Worked for the CIO of Vision and Strategy in a team that reimaged SSA architecture and software from the ground up using modern technologies (Java/J2EE + HBase) to improve experience and save \$1B a year
- Designed high-level strategy & data model and implemented prototype service-oriented architecture and IaaS
- Developed RESTful eligibility services using Drools rules engine and a “Turbo-Tax”-like, Struts2 web frontend

Education & Skills

Cornell University, Ithaca, NY

Master of Engineering in Computer Science, May 2010

Bachelor of Arts in Government and Computer Science, May 2009

Relevant Courses: Distributed Systems, Algorithms, Machine Learning, Artificial Intelligence, Algorithms, Scientific Computing, Security

Exchange Student (2008) at **Freie Universität Berlin, Berlin, Germany** (functionally fluent in German)

Programming Languages: *Experienced:* Python, Ruby (+Rails & Sinatra), Java/J2EE, SQL, HTML5/CSS 3/JavaScript;

Comfortable: Scala, Node.js, PHP; *Familiar:* Go, Objective C, C/C++, SML and MIPS Assembly

Other: Amazon Web Services, Google Cloud Platform, Terraform, DynamoDB, Redis, Airflow, Spark, MapReduce, PostgreSQL, MySQL, Linux, Bash, Git, LDAP, Photoshop/Illustrator, Microsoft Office