Xuanqi Li

Sunnyvale, CA 94089 | 408.406.8721 | lixuanqi1995@gmail.com
github.com/LiXuanqi | www.linkedin.com/in/xuanqi-li | xuanqi.dev

EDUCATION

Northeastern University **09/2018 – 12/2019**

- M.S. Information Systems, concentrating in Computer Software Engineering, GPA 3.89

WORK EXPERIENCE

Airtable, Senior Software Engineer **04/2024 – Present**

San Francisco, CA

- Airtable is a no-code platform for building custom databases and workflows with a spreadsheet-like interface. My team is responsible for improving application **vertical scalability**, a critical initiative that grew from **6 to 14 engineers** since I joined.
- Core contributor to in-house **relational database built from scratch**, implementing cascading deletion system for removeModel operations (DROP TABLE with FK/relationship cleanup), consolidating **Rust-TypeScript interop** protocols for **30% improved data exchange performance**, and optimizing cold-load performance through deferred index statistics
- Orchestrated seamless migration from legacy in-memory systems to new relational database architecture, designing **dual-write strategy** to gradually migrate data and business logic without service disruption
- Led multi-quarter re-architecture of **Node.js** request processing framework from **single-threaded to multi-threaded** implementation, delivering significant throughput improvements for high-concurrency workloads
- Enhanced team productivity by configuring **Rust-analyzer** integration with custom **Bazel** build system, enabling full IDE support (auto-completion, linting, error detection) for Rust development

LinkedIn, Senior Software Engineer **08/2021 – 04/2024**

Sunnyvale, CA

- Owned end-to-end backend development for **LinkedIn Pages** admin experience, empowering **67M+ businesses** to manage company presence and analyze follower engagement
- Collaborated with cross-functional teams (PM, Design, Marketing, Support) and platform teams (Premium, Feed, Invitations, Infrastructure) to deliver weekly analytics notifications, page-to-page following, invitation credits, competitor analytics, and dashboard redesign
- Built **Java** APIs and **Spark** data pipelines integrating with internal **NoSQL databases** (document, graph, time-series) for real-time insights and batch processing
- Pioneered **Server Driven UI (SDUI)** architecture adoption, reducing feature development time by **67%** and eliminating mobile app releases for UI changes
- Productivity Working Group initiatives - conducted engineer interviews, leveraged **LLM analysis** to prioritize improvements, co-authored company-wide best practices, and championed modern tooling adoption (**JUnit, AssertJ, GitHub Copilot**)
- Mentored interns and frontend engineers transitioning to backend, supported quarterly OKR planning, and recognized with **MVP award** and **Top Interviewer** distinction

Yelp, Software Engineer at Ads Platform **01/2020 – 08/2021**

San Francisco, CA

- One of the developers/maintainers for the **Yelp Campaign Management APIs**.
- Developed the customization features (e.g. Location targeting, Keywords targeting) for campaigns with **AWS SQS**, **Kafka**, **StreamSQL**, **Redshift**
- Migrated daily ad inventory ETL from **MapReduce** to **Spark** and reduced the **75%** runtime and cost.
- Led the project to integrate our experimentation platform into Ads team, cooperated with stakeholders across teams, and advocate new tool in Ads team (~50 engineers).
- **Release engineer** for deploying code across multiple teams
- **Mentored new hires** and helped them onboard smoothly and quickly.
- Joined **on-call** rotation to debug production issues across Ad system.
- Became the on-point who answered any question related to our system in team slack channel.
- One of the **instructors of Ads academy** and teached ~ 20 new hires each cohort.

Yelp, Software Engineering Intern **05/2019 – 08/2019**

San Francisco, CA

- Designed the schema of new **Cassandra** table to support more complex and flexible query.
- Modified a **MapReduce** batch and persisted data into **AWS S3**, **Cassandra** and **Redshift**
- Created a **RESTful API** using **Pyramid** and manipulated **MySQL** tables with **SQLALchemy**.
- Built a **Python** batch for comparing or backfilling Cassandra tables in the past two years(140GB) and solved the performance problems of querying without partition key.