

Yaming Zhang

Mountain View, CA

zhangyaming0726@gmail.com

yamingzhang.com

857-437-2701

EDUCATION

Northeastern University, Boston, MA

Master of Science in Computer Engineering

Dec 2023

University of California Riverside, Riverside, CA

Graduate Preparation Program in Computer Science

Jun 2020

Xidian University, Xi'an, China

Bachelor of Science in Computer Science

Jun 2019

WORK EXPERIENCE

Moveworks (now part of ServiceNow), Mountain View, CA

Senior Software Engineer, Core Infrastructure

Jan 2024 – Present

- Designed and implemented **funnel-tracing** across the Minerva AI chat flow; primary data source for ML, product, and applied-research teams measuring AI assistant quality.
- Designed and built Moveworks' **Apache Spark Structured Streaming** platform from scratch (Kafka extractor, streaming loader, K8s job submitter, observability) — backbone for real-time ML pipelines.
- Primary owner of Moveflow, Moveworks' **Apache Airflow** orchestration platform for all batch/ETL jobs; built the PagerDuty alerting framework, shipped OpenTelemetry metrics across all prod regions, and hardened the scheduler.
- Re-architected Recon config-management: solo-replaced go-git with native git for **~8–9× faster startup** (16 Go deps eliminated); built an async Kafka change-event pipeline.

Electronic Arts (EA), Redwood City, CA

Software Engineer Intern

May 2023 – Aug 2023

- Designed the **Java-based OOP** model for Player Lifecycle Graph, a serialization model to manage players' states and unified messages, resulting in a **30% cost reduction** by eliminating the need for customizing **Braze**.
- Developed a **Flink** topology for **Kafka** streaming PIN events execution, enabling efficient retrieval of lifecycles from **Redis**, reading player profiles from **ScyllaDB**, and executing the corresponding lifecycles.
- Leveraged EADP's **Data Query Language (DQL)** to extract data from unstructured **JSON** events, evaluate boolean conditions, and execute action scripts.
- Built a web application with **React** and implemented APIs with **Spring Boot**, empowering campaign managers to draw lifecycle flowcharts and retrieve players' current states in games.

Tencent, Beijing, China

Software Engineer

Jul 2020 – Aug 2022

- Developed **Microservices** on **Go** using **gRPC** with logs on **ELK**, monitoring on **Prometheus**, asynchronous pipelines to **Kafka**, etc, achieving a **25% improvement** in interface response time and a **50% reduction** in production debug time.
- Established automated **CI/CD** pipelines that enhanced service security by seamlessly integrating **Unit Tests** during CI, while also enabling efficient rollback on **Tencent Cloud Kubernetes**.
- Wrote customized codes in the secondary development of **Superset** to create a dashboard displaying user usage data of Tencent services by preprocessing and visualizing data with **Pandas** and **ECharts**.
- Analyzed the traffic data of Tencent news using **Scala Spark** and **PySpark** to promote automatic advertising and personalized recommendation.

Big Data Lab, UC Riverside, Riverside, CA

Research Assistant — Advisor: Prof. Ahmed Eldawy

Aug 2019 – Jul 2020

- Designed and benchmarked computational-geometry primitives (WKT/WKB conversions) across **Geolite**, **Esri**, and **JTS** libraries on Apache Spark; published as **first author** at the **ACM SIGMOD GeoRich Workshop, 2020** (11 citations).
 - Built the OpenStreetMap extraction pipeline for **Beast** — the lab's spatio-temporal analytics system — using **Scala**, **Spark**, and the **Hadoop** file API to parse PBF data on **AWS EMR**; foundation for the **OSMX** paper at **ACM SIGSPATIAL, 2022**.
 - Optimized Beast Raptor spatial-join performance by increasing join granularity from pixels to tiles, enabling large-scale satellite-image-over-OSM-vector joins for downstream spatial ML.
 - Strengthened the **UCR-STAR** I/O module by implementing KML output and compression in Java; co-author on **Beast: Scalable Exploratory Analytics on Spatio-temporal Data** at **ACM CIKM, 2021** (43 citations).
-

PUBLICATIONS

Samriddhi Singla, **Yaming Zhang**, Ahmed Eldawy (2022). [OSMX: Spark-based Geospatial Data Extractor from OpenStreetMap and Satellite Images](#). International Conference on Advances in Geographic Information Systems.

Ahmed Eldawy, ..., **Yaming Zhang** (2021). [Beast: Scalable Exploratory Analytics on Spatio-temporal Data](#). ACM International Conference on Information and Knowledge Management.

Yaming Zhang, Ahmed Eldawy (2020). [Evaluating Computational Geometry Libraries for Big Spatial Data Exploration](#). International ACM SIGMOD Workshop on Managing and Mining Enriched Geo-Spatial Data.