

Albert (Geyang) Xu

858-568-6771 | albertxu1010@gmail.com | [LinkedIn](#) | [Portfolio](#) | San Diego, CA

EDUCATION

University of California, San Diego

Sep 2022 - Mar 2024

M.S in Computer Science (GPA: 3.88/4.0)

La Jolla, CA

- Courses: Principles of Computer Operating Systems, Principles/Program Languages, Statistical NLP, Principles of Computer Architecture, Computer Science and Engineering, Principles of Machine Learning, Principles of AI

University of Liverpool

Sep 2018 - Jun 2022

B.S in Computer Science (GPA: 3.93/4.0, top 3%)

Liverpool, UK

TECHNICAL SKILLS

Programming Languages: Python, Java, Go, SQL, C/C++, HTML/CSS, JavaScript, Rust, Objective-C, C#, R, Haskell

Frameworks and Tools: Django, React, Git, Docker, AWS, gRPC, Matlab, Flask, Hadoop, DynamoDB

EXPERIENCE

4Pexonic Inc.

Oct 2024 - Present

Back-end Engineer

San Diego, CA

- Built a one-stop research-to-publication platform on a **Go** micro-service stack with **gRPC/REST** and **NATS** JetStream; webhook fan-out decoupled services and achieved **3x peak throughput**
- Replaced fuzzy SQL search with **ElasticSearch** full-text indexing, cutting query latency by **30%**
- Constructed multi-stage **Docker** builds and GitHub Actions blue-green **CI/CD** pipelines, achieving zero-downtime releases, **50%** faster deploys, and automated rollbacks in under five minutes
- Added a **Redis** read-through cache with batched writes, cutting p95 latency **40%** and **2x QPS**

UCSD ECE Department

Feb 2024 - Jun 2024

Web Developer

La Jolla, CA

- Developed and maintained the ITA Workshop [website](#), leveraging **Django** and **React**, serving **1000+ users**; added **Redis** server-side caching and targeted **SQL indexes**, trimming average page-load time by **35%**
- Refactored video / document uploads to **asynchronous** chunked processing with **UUID** file names, boosting throughput **20%** and preserving **ACID**-safe schema changes
- Migrated large-file uploads to **multipart S3 transfers** orchestrated by **Celery** with resumable retries and SHA-256 integrity checks, sustaining **1GBps** peaks without data loss
- Managed **PostgreSQL** and configured **Amazon S3 cross-region replication** for versioned backup

UCSD Halicioğlu Data Science Institute

Oct 2023 - Sep 2024

Research Engineer

La Jolla, CA

- Built a modular “Injector” pipeline (pattern-gen → sampling → injection → evaluation) with **beam-search** pruning and Optuna Bayesian tuning, raising downstream model accuracy by **10%** and fairness by **15%**
- Designed and executed a full Inject → Clean → Retrain benchmark to stress-test cleaning tools under missing-value, selection-bias, and outlier scenarios; results accepted for **VLDB** revision
- Improved the standardization method for injecting errors (**Z-score** normalization) to overcome uneven data distribution, enhancing the framework’s attack capability by **30%**

PROJECTS

GovChat: Smart-City Conversational Platform

- Deployed **GPU**-boosted **LLM** micro-services for domain-specific responses with **35%** lower latency under **200 ms**
- Engineered a three-tier cache—in-process **LRU**, **Redis Sorted Set** (pagination < 1ms), and **MongoDB aggregation**—offloading **80 %** of read traffic and eliminating deep-paging stalls
- Built a **Kafka** streaming pipeline for decoupling alert ingestion, maintaining high throughput with **p99 ≤ 50 ms**
- Hardened security with **bcrypt** and **Nginx** TLS 1.3, achieving **99.9%** uptime and **40%** fewer unauthorized requests

SurfStore: Distributed Cloud Storage Server Project | [Link](#)

- Designed and implemented a **Dropbox**-like cloud-based file storage service using **Go**, featuring distributed storage and consistent hashing algorithms for efficient file management
- Built for large-scale data handling, the project showcased excellent scalability during tests and efficiently met the data storage demands of thousands of users through client-server communication via **gRPC**
- Enhanced the system's fault tolerance by implementing the **RAFT** distributed consensus protocol, which led to a **50%** improvement in error recovery efficiency during testing