

# ASHRAY CHOWDHRY

ashraychowdhry@gmail.com | linkedin.com/in/ashraychowdhry | ashraychowdhry.com | (609) 937-9394

## EXPERIENCE

### RIPPLE LABS | *Software Engineer II, I*

Aug 2022 – Present | New York, NY

- Tech lead of a 6-person team, developing open-source XRPL libraries (xrpl.js, xrpl.py, xrpl4j, xrpl-rust) and new protocol amendments, triaging community issues, and shipping tooling improvements used across the XRPL developer ecosystem.
- Head of a 4-person team that architected, developed, and shipped los.prod.ripple.io, a public data indexer consolidating complex market blockchain data for XRPL into unified API endpoints for low latency processing. Serves as the official ETL service that powers analytics for 15,000+ daily users, processing up to 100,000 tx/min.
- Directed delivery of AMM, Vault, Lending Protocol, MPT, Tokens and real-time search capabilities, enabling high-fidelity queries over DEX trades, order books, and market state not available in any prior public XRPL service.
- Advocated for and implemented a re-architecture for Ripple's OpenSearch indexing to cut p95 latency from ~15s to <2s, enabling real-time interactive exploration of on-chain data that was previously impractical for developers to query.
- Founding engineer for Ripple's institutional onboarding portal, improving customer time to first payment and risk analysis reliability by over 60%. The service has since grown into its own organization at Ripple.
- Built a queue-backed audit trail microservice for Ripple Payments to meet growing compliance requirements, along with contributions to smart liquidation service, and payments routing. The platform expanded cross-border support from 20 to 50 banks globally.

### MAKESWIFT | *Consultant Software Engineer*

Apr 2022 – Aug 2022 | Atlanta, GA

- Created automated component generation tooling in Bash, GraphQL, TypeScript, React, for a no-code platform, producing thousands of CMS-driven components and enabling rapid client migrations including BigCommerce.

### CAPITAL ONE | *Software Engineer*

May 2021 – Sep 2021 | Richmond, VA

- Designed and built a full-stack customer-facing retirement planning product with React, Flask, PostgreSQL, and AWS.

### JOHN DEERE FINANCIAL | *Data Analytics Engineer*

May 2020 – Sep 2020 | Des Moines, IA

- Developed and productionized a customer churn predictive model to improve buying trend analysis based on 40-factor raw data with PySpark, SQL, Databricks, and MLflow. Yielded a 35% increase in benchmark performance and customer retention spend.

### GT SUSTAINABILITY LAB | *Systems Eng. Modeling Researcher*

Dec 2019 – Jun 2022 | Atlanta, GA

- Co-authored a published IEEE Systems Journal paper on a novel systems metric quantifying resilience, derived from ecological systems to improve human-made networks.
- Spearheaded a modelling tool to simulate large-scale ecosystems and stress-test human-made systems with MATLAB and Node.js

## EDUCATION

### GEORGIA INSTITUTE OF TECHNOLOGY | Atlanta, GA

B.S. in Computer Science, Certificate in Finance | 3.95 GPA

Relevant Courses: NLP, Machine Learning, Networking, Database Design, Data Structures, Fin. Modeling, FinTech/Cryptocurrency

## PROJECTS

### SEATBELT | *AI-Powered Autonomous E2E Testing Agent*

- Multi-agent AI testing platform that explores web apps like a first-time user, discovering bugs, security vulnerabilities, and UX issues. Built with autonomous LLM orchestration, multimodal vision, structured tool use, perceive-decide-act loops, parallel execution with shared state, RAG knowledge base, context window optimization, CI/CD integration, and three-phase AI pipeline architecture.

### DESIGN DNA | *UI/UX Intelligence Extraction Agent Engine*

- Compound agentic AI tool that reverse-engineers design systems from live websites— capturing CSS animations, pseudo-state styles, scroll-triggered motion, and responsive layouts and synthesizes Tailwind configs. Uses multimodal prompt chaining, vision-based agentic self-verification loops, adaptive model routing with cost-aware depth scaling, parsing with fallback chains, cross-call data fusion to prevent hallucination, and role persona engineering across an Analyst-Generator-Critic pipeline.

### TROVE AI | *AI-Powered Content Memory Platform — iOS App Store*

- Full-stack iOS app for saving content across various social media platforms. Uses on-device NLP, LLM enrichment, and Playwright extraction, RAG-powered conversational search over user content with hybrid vector + keyword retrieval, real-time cross-user trending analytics with fuzzy deduplication, parallel server-side LLM orchestration with graceful degradation, and a share extension processing queue — achieving sub-50ms time-to-first-content and 91% extraction success.

## SKILLS

**Languages:** Python, C++, Rust, Java, TypeScript/JavaScript, C, SQL, HTML/CSS, Swift, Solidity, MATLAB

**Infrastructure:** AWS, GCP, Terraform, Docker, Kubernetes, Supabase, Cloudflare, Vercel, Prometheus, Grafana, Git, Argo, Kafka

**Databases:** PostgreSQL, MySQL, MongoDB, CockroachDB, PgVector, OpenSearch, DynamoDB, S3

**Frameworks:** PyTorch, Tensorflow, MLflow, LangChain/LangGraph, Spark, React., Next.js, Node.js, Flask, Spring Boot, Arduino