

MJ Moshiri

415-218-4540 · thisismjmoshiri@gmail.com mjmoshiri.com github.com/mjmoshiri ·

Senior Software Engineer building **production agentic systems**, and scalable full-stack solutions. Proven track record of delivering significant cost reductions, latency improvements, and operational scale across YC startups and large enterprises. Expertise in translating complex product, fraud, and compliance requirements into reliable, **high-performance architecture**.

EXPERIENCE

Senior Software Engineer – Oway (YC S24)

Dec 2025 – Present

- Built **core infra for agentic workflows**, including execution patterns, tools, policies, and internal service integrations.
- Developed email workflow infrastructure for shippers, carriers, and internal operational teams.
- Established **engineering standards** across architecture, testing, code quality, and AI-assisted development.
- Designed the foundation for coverage workflows to enable future data-driven and ML-based decision making.
- Self-hosted latency-critical services, **significantly reducing costs** and improving control over data and compliance.

Software Engineer – Visa Inc

May 2025 – Dec 2025

- **Led development** of an LLM-powered compliance system for Visa+ user flows against policies and brand standards.
- Re-architected the platform, reducing p99 review latency from **35 minutes to 3 minutes**.
- Authored design documents on **RAG quality, evaluation frameworks**, and human-in-the-loop AI systems.
- Trained subject-matter experts to translate policy knowledge into reliable AI-assisted workflows.
- Led LLM onboarding for an eight-engineer team, covering prompt design and GenAI best practices.
- Owned **GenAI demos** presented at Visa's **global executive leadership** market event.

Member of Technical Staff – SafetyKit (YC S23)

Dec 2024 – May 2025

- Led forward-deployed engineering for enterprise customers, translating compliance requirements into AI workflows.
- Helped scale revenue from mid-six figures to a **multimillion-dollar run rate**.
- Re-architected an agentic pipeline, **reducing p99 cost from \$0.20 to \$0.015 per request**, while keeping accuracy.
- Built **multimodal fraud detection systems** using model cascades, SAM2, CLIP, and image-forensics techniques.
- Built infra for agentic computer-use workflows, integrating LLMs, browser automation, AWS, and internal services.
- Designed prompts, policies, and **evaluation frameworks** across 20+ safety and compliance categories.

Software Engineer – Freelance

Apr 2024 – Dec 2024

- Contracted with xAI on synthetic STEM data for Grok training workflows.
- Designed RAG architecture for YC startup: **retrieval precision 63% → 87%** across 10K+ documents.

Software Engineer Intern – Nokia

Dec 2022 – May 2023

- Improved API **throughput 20% and latency ~15%** via caching and code-path optimization.
- Led Node.js 14→18 and TypeScript 3→5 migration for a product deployed across 100+ countries.

SKILLS

Languages: Java, TypeScript, Python, Go, SQL

AI/ML: LLMs, RAG, Context Engineering, LangChain, Transformers, Evaluation, Model Cascades, Guardrails

Data: PostgreSQL, Cassandra, MongoDB, BigQuery, Weaviate, Spark, InfluxDB

Infrastructure: AWS, Docker, Kubernetes, OpenShift, Linux, Distributed Systems

Tools: Git, Bash, FastAPI, React/Next.js, REST APIs, Microservices