

Michael Lubich

San Francisco, CA | +1 (415) 275-0094 | michaella.lubich@gmail.com | [LinkedIn](#) | [Github](#) | mishalubich.com
U.S. Citizen (W-2) | Authorized to work in the U.S. without sponsorship | **Job preferences:** Onsite, Hybrid, or Remote

PROFESSIONAL SUMMARY

Backend software engineer with **ex-Apple** and **ex-Walmart** background shipping production systems and high-throughput microservices – **Java/Spring Boot**, **Python**, and **TypeScript**. **Tripled microservice throughput** at Walmart and validated **macOS APFS storage releases across 100M+ devices** at Apple. Built and operates a support-triage agent handling **10,000 cases/month** at a **zero hallucinated-write rate**. Ships REST/gRPC services, **Kafka/ClickHouse** data paths (**sub-200ms** latency), and cloud/CI on **AWS** and **Kubernetes**.

SKILLS

Languages	Java, Python, TypeScript, JavaScript, SQL, Go, Bash, C/C++
Backend	Java/Spring Boot, Spring MVC, Hibernate, JPA; FastAPI, Flask, Django; Node.js, Express; REST, GraphQL, gRPC; microservices, serverless (Lambda), event-driven (Kafka, RabbitMQ); SOLID, DDD, design patterns
Data	Postgres, Oracle, MySQL, Redis, MongoDB; Snowflake, BigQuery, Elasticsearch; ETL pipelines, Spark, Airflow, Kafka; Pandas, NumPy
Cloud	AWS (EC2, ECS, Lambda, S3, RDS, DynamoDB), Azure (App Services, DevOps), GCP (GKE, BigQuery); Terraform, Ansible, Docker, Kubernetes, Helm
DevOps	Git, GitHub, GitLab, Jenkins, GitHub Actions, Bitbucket Pipelines; Maven, Gradle; CI/CD pipelines, infrastructure as code, GitOps, runbooks
Testing	JUnit, Mockito, TestNG; PyTest, Playwright, Selenium; Locust, k6, JMeter; CI gates, flaky-test triage, TDD
Observability	Prometheus, Grafana, OpenTelemetry, Datadog, ELK, Sentry, Splunk, New Relic
Security	Enterprise cybersecurity; OAuth 2.0, SSO/SAML, RBAC, IAM; OWASP Top 10, OWASP LLM Top 10; prompt injection defense, tool-abuse controls, least-privilege design
Frontend	React, Next.js, Angular, Tailwind; Figma, Postman
AI / ML	On-prem agents; LangChain, LangGraph, agentic workflows, agent loops and harnesses (Claude Code, Hermes, OpenClaw); MLOps and ML pipelines, LLM fine-tuning (LoRA/QLoRA); open-weight models (GPT-OSS 120B, Llama, Qwen, DeepSeek) via vLLM; MCP tool servers, agent skills, multi-app synthesis; deterministic RAG for hallucination control; frontier APIs (OpenAI, Anthropic Claude, Bedrock); prompt and context engineering, LLM evaluation, PyTorch, scikit-learn
Methodologies	Agile/Scrum, TDD, code review, mentorship, on-call, software delivery

EDUCATION

- University of California, Berkeley** | *Bachelor of Arts, Computer Science; GPA 3.7*

Berkeley, CA | Aug 2019 – Aug 2023

PROFESSIONAL EXPERIENCE

- EchoStar Corporation** | *Staff AI Engineer* North America (Hybrid) | Sep 2026 – Present
 - Multi-Agent Systems for Real-Time Telemetry Triage** Orchestrates **multi-agent AI systems** solving real-time telemetry triage, incident diagnosis, and root-cause analysis across distributed network infrastructure.
 - Distributed Agent Orchestration Frameworks** Architects distributed agent orchestration frameworks and **deterministic state graphs** managing task delegation and state retention across enterprise APIs.
 - Runtime Guardrails Against Non-Deterministic Failure** Engineers runtime execution guardrails, dynamic **context management**, and secure **tool-calling** protocols preventing non-deterministic agent failures.
 - Domain Agents in Production Telemetry Pipelines** Integrates domain-specialized agents into production **telemetry data pipelines**, automating end-to-end diagnostic and operational workflows.
 - Standards and Benchmarks for Enterprise Agent Deployments** Defines organizational architectural standards, **evaluation harnesses**, and reliability benchmarks for enterprise-wide agent deployments.
- Independent Consulting (W3 Sourcing, ERIA, LUPFR)** | *Technical Partner & Software Engineer* San Francisco, CA | Mar 2026 – Present
 - Production Backend for Live Case Triage** Shipped a **Python** case-triage service into a live Salesforce org under a six-field write allowlist and ~30 evidence gates, working **10,000 cases per month** at a **zero hallucinated-write rate** and returning 40 hours per week to the support team.
 - Full-Stack Commerce Platform (ERIA)** Technical partner for a \$10M+/yr events business serving Anthropic, OpenAI, and Johnson & Johnson, building a **Next.js** and **Supabase** commerce platform on **Stripe** with webhook idempotency and refund handling, shipped against **2,895 tests**; separately automated operator self-service for a non-technical events team (lupfr.com), replacing recurring manual publishing work. (www.eria.co)
 - Agentic ATS Backend (W3 Sourcing)** Built a Next.js portal and Python CLI over **19K+ candidate records** with schema-gated scoring output and an **85% coverage gate**, backing a recruiter-facing sourcing and screening funnel. Replaced **\$20K/yr** of vendor tooling; the firm projects **\$500K–\$1M/yr** in placements off the pipeline. (www.w3sourcing.com)
- Polaris Wireless** | *Platform & AI Software Engineer* Santa Clara, CA | Sep 2024 – Jul 2026
 - Agentic Data Engine (AigisQuery)** Owns the plain-English query platform non-technical teams use against carrier geospatial data, built on **Model Context Protocol (MCP)** tool servers under least-privilege execution.
 - Carrier-Scale Ingestion (Kafka / ClickHouse)** Scaled **Kafka**, **ClickHouse**, and **ElasticSearch** pipelines over billions of real-time spatial records, holding query latency **under 200ms** in production.
 - Agentic Harness and Execution Guardrails** Built the production-grade agentic harness every agent runs on – deterministic loop detection, step budgets, and write restrictions.
- Apple** | *Software Development Engineer in Test* Cupertino, CA | Jan 2023 – Jun 2024
 - Automated Validation Pipelines** Owns modular **macOS APFS** validation pipelines that hold releases stable across **100M+ devices**, cutting manual maintenance and accelerating deployment cycles.
 - Production Defect Resolution** Triageed and resolved **1,100+ high-priority production defects** through structured root-cause workflows, keeping APFS releases stable across consumer and enterprise environments.
 - Deployment Workflows** Built **Ansible** and **Docker** deployment workflows that removed manual configuration steps from the release path and accelerated release cycles.
- Walmart** | *Software Engineer Intern* Sunnyvale, CA | May 2022 – Aug 2022
 - High-Throughput Backends** Shipped backend services that **increased throughput 300%** using **Java**, **Spring Boot**, and **AWS RDS**, deployed as horizontally scalable microservices on Kubernetes with Kafka.
- Lawrence Berkeley National Laboratory** | *Software Engineer* Berkeley, CA | May 2021 – Sep 2021
 - Containerized ML Pipeline** Replaced manual experimentation with a versioned pipeline (**Python**, **Docker**, **Airflow**, GridSearchCV, Optuna, MLflow), **saving 200+ hours annually**.
- Honda Innovations** | *Software Engineer Intern* Mountain View, CA | Jan 2021 – May 2021
 - Route Optimization and Deployment** Built a predictive route-optimization model for medical-supply delivery and shipped it to **AWS ECS** behind an automated **Docker/Terraform** CI/CD pipeline.

PROJECTS

- Event-Driven Order Service** **Java/Spring Boot** order-service prototype with **Kafka** event sourcing, Postgres write model, Redis cache, and Terraform-provisioned AWS deployment; gRPC + REST surfaces with JUnit/Mockito coverage.
- imail-mcp – Autonomous Email Agent** Built an Apple Mail agent (**Python**, AppleScript, MCP tool server) that reads, searches, and organizes a live inbox and drafts replies, turning hours of daily manual triage into a surfaced action list – with sends gated behind explicit per-account targeting so the agent can never mail the wrong list. Published open source. ([imail-mcp](#))

PUBLICATIONS

6 peer-reviewed papers in machine learning and AI for environmental science ([Google Scholar](#)): **1. Stream temperature predictions using ML – Water (MDPI), 2022** **2. Classical ML for stream temperature predictions – AGU Fall Meeting, 2022** **3. Climate-driven disturbances on river water quality – Frontiers in Hydrology, 2022** **4. ML/statistical modeling for river water quality – AGU Fall Meeting, 2021** **5. Data-model integration for hydrobiogeochemical modeling – AGU Fall Meeting, 2021** **6. Predicting stream temperature with low-complexity ML – AGU Fall Meeting, 2021**