

# O. FELIX AMORUWA

famoruwa@berkeley.edu | 909-731-9011 | [felixamoruwa.info](mailto:felixamoruwa.info)

Staff Product Manager with 12+ years owning data platforms, cloud infrastructure, and API-first services at scale — from distributed search and PostgreSQL metadata catalogs at Splunk to a 675M-engagement platform running at 50K TPS with sub-25ms latency at Intuit. Deep technical fluency across relational/distributed databases, streaming pipelines, and multi-cloud architectures, with hands-on experience building vector search, real-time analytics, and LLM serving pipelines for next-generation AI workloads. Track record delivering 0-to-1 data services and driving 1-to-N scale in fast-moving, high-growth environments. NeurIPS published; UC Berkeley Engineering.

## PROFESSIONAL EXPERIENCE

---

### SPLUNK INC. | San Francisco, CA

*Senior Product Manager — Search Orchestration & Data Services* 01/19 – 01/21

- Owned Search Catalog (PostgreSQL metadata service), Search Service (Go microservices), and Splunk Processing Language (SPL/SPL2) — managing the full data services portfolio from metadata and catalog to query execution for Splunk Cloud Services (SCS).
- Delivered Scheduler Service end-to-end in ~4 months (June–Oct 2019), enabling scheduled search and real-time data pipeline capabilities for first-party applications; demoed at Splunk .conf19.
- Led query performance optimization initiative for beta customer (Assurance), building mirrored Enterprise topology for benchmark testing and achieving up to 10x throughput improvements in SCS search — directly addressing scale, throughput, and performance for cloud-scale data workloads.
- Designed repeatable RICE-based prioritization framework across 3 microservice backlogs, balancing internal partner, third-party developer, and Fortune 500 customer requirements. Splunk .conf18 Speaker.

### INTUIT | Mountain View, CA

*Staff Product Manager — Developer Frameworks & Platform Infrastructure* 06/21 – 09/24

- Scaled ICE platform to 675M+ engagements in FY23 (275% YoY growth) across QuickBooks, TurboTax, Mint, Mailchimp, and Credit Karma; drove throughput from 6K to 50K TPS via rSocket migration supporting ~1.5M concurrent connections with sub-25ms TP99 — demonstrating the scale and performance required for next-generation AI cloud workloads.
- Built Asterias, a declarative asset lifecycle management platform with GraphQL API, providing metadata catalog and governance capabilities across ~20 mobile apps and 30+ product SKUs; used SQL and BigQuery telemetry to prioritize developer pain points at enterprise scale.
- Led Mailchimp GCP-to-AWS migration for MSaaS, delivering Golang service template, MySQL persistence integration, and updated DevPortal documentation — hands-on multi-cloud architecture experience spanning data locality, replication, and production cutover.
- Initiated MSaaS Drift Detection and Resolution program: authored Java JAR library to scan Git repos for configuration drift and built remediation roadmap using OpenRewrite — data governance and catalog integrity at enterprise scale.
- Delivered ICE Self-Service platform (DevPortal, GitOps config, ICE Playground), reducing developer onboarding from 2–3 weeks to minutes in pre-prod and <24 hours for production, while mitigating \$1M+ in projected opex growth.
- Implemented ICE Presence in async chat generating \$480K/month in additional invoicing; deployed Background-to-Foreground Messaging on iOS/Android with <100ms latency. DeveloperWeek 2022 speaker.

### FINTELLECT AI | Oakland, CA

*Founder & CEO — fintellectai.io (iOS/Android/Web)* 09/24 – Present

- Architected RAG retrieval pipeline with ChromaDB vector store, multi-provider LLM orchestration (Claude, GPT-4, Gemini) with fallback routing, structured output validation, and token budget optimization — building AI/ML data products including vector search and LLM serving pipelines for real-time financial analytics.
- Built AI-powered real-time analytics and macroeconomic charting tools integrating LLM models for automated trade analysis; executed go-to-market through App Store launch with tiered subscription pricing.

### STREAMIO AI | Oakland, CA

*Founder & CEO — streamio.ai (macOS, Linux, iOS & Web)* 09/24 – Present

- Engineered real-time HLS livestreaming pipeline with multi-stream canvas compositing (up to 9 concurrent sources at 30fps), FFmpeg transcoding, WebSocket communication layer, and local HTTP playback server — end-to-end streaming and data pipeline architecture from ingest to serve.
- Implemented OpenClaw multi-agent orchestration framework with gateway protocol and subagent delegation; built end-to-end auth and payments pipeline with Kinde OAuth 2.0, Stripe tiered subscriptions, and Electron SafeStorage for secure credential management — usage-based pricing and data security in production.
- Shipped cross-platform desktop builds with native macOS ScreenCaptureKit integration, embedded terminal (20+ slash commands), and MCP server exposing screen capture tools to AI coding assistants — API-first platform design serving as critical infrastructure for AI workflows.

### KAISER PERMANENTE | Pleasanton, CA

*SOA Technical Product Manager* 08/12 – 12/18

- Led development and enterprise rollout of Splunk Logging-as-a-Service (1.7 TB daily ingest volume, 200+ internal enterprise customers) and ITSI Application Monitoring-as-a-Service — managed data services at enterprise scale.
- Built distributed caching capability using Redis and XC10 across the enterprise to address scalability, fault tolerance, and data redundancy; led demand forecasting and capacity planning for SOA product suite across multiple datacenters.

### IBM | Pittsburgh, PA / Foster City, CA / RTP, NC

*Software Engineer — Business Intelligence Products* 07/05 – 03/10

- Resolved multiple \$1–5M high-severity customer escalations for enterprise BI and data platform products; led cross-functional root cause analysis and customer-facing resolution, improving case resolution time by 20% across the Support Organization.

### BANK OF AMERICA MERRILL LYNCH | Charlotte, NC

*Tech MBA Summer Associate — Global Corporate & Investment Banking T&O* 06/11 – 08/11

- Developed enhanced estimation model using Monte Carlo simulation to optimize DMAIC phase distribution across a \$494M portfolio in Global Technology & Operations.

## AI/ML RESEARCH & PROJECTS

---

### aeval — AI Model Evaluation Platform (2025–2026)

- Built local-first evaluation platform with FastAPI orchestrator, TimescaleDB (time-series database), Redis job queue, and Next.js dashboard — API-first data platform architecture with managed database, streaming job queue, and real-time analytics for 5 eval types (factuality, reasoning, instruction-following, safety, code generation).

- Implemented data governance and quality gates: CI/CD regression detection, automated safety gates, data contamination detection via SHA-256 hashing, and statistical rigor (bootstrap confidence intervals, Welch's t-test, Cohen's d effect size).

#### **RL Workbench — Post-Training RL Platform** (2026)

- Built 3-phase RL post-training workbench with live SSE metric streaming, GPU Docker passthrough for framework benchmarking (TRL, VeRL, OpenRLHF, NeMo RL) on CUDA/Apple Silicon — firsthand experience with the GPU-intensive AI workloads CoreWeave's customers run at scale.
- Implemented 12 RL algorithms (PPO, GRPO, DAPO, DPO, SimPO, and 7 others) with standardized throughput/memory/convergence benchmarking and cross-tab workflow lineage tracking across 4 datasets (GSM8K, MATH, HumanEval, UltraFeedback).

#### **BRAIN — Protein Structure Prediction ML Platform** (UC Berkeley 2004; rewritten 2026)

- Built production ML platform in PyTorch with MLflow experiment tracking, Optuna HPO, FastAPI serving, and Docker orchestration (6 containers); NeurIPS 2014 accepted paper. Original 2004 system: hand-coded neural network in C++ with custom BPTT; 2026 rewrite spans 413 to 8B parameters.

#### **AutoEval — Automated Visual Evaluation for Robot Model Training** (2025)

- Built zero-integration automated evaluation pipeline capturing from any visualization tool (RViz, Matplotlib) via screen capture; multimodal AI (Claude/GPT-4V) performs spatial reasoning on prediction frames, generating structured PASS/FAIL reports — reducing evaluation cycles from 72 hours to ~4 minutes.

## **EDUCATION**

---

### **CARNEGIE MELLON UNIVERSITY, SILICON VALLEY**

Master of Science — Software Management Mountain View, CA (08/16)

### **CARNEGIE MELLON, TEPPER SCHOOL OF BUSINESS**

Master of Business Administration — Concentrations: Finance, Quantitative Analysis, Economics Pittsburgh, PA (03/12)

### **UC BERKELEY, COLLEGE OF ENGINEERING & HAAS SCHOOL OF BUSINESS**

Bachelor of Science — Computational Engineering Science Berkeley, CA (12/05)

Certificate — Corporate Finance for Financial Engineering

## **TEACHING EXPERIENCE**

---

### **DE ANZA COLLEGE** | Cupertino, CA

*Adjunct Faculty — Computer Information Systems Department* 04/18 – Present

- CIS 35A/B, 36A/B — Java Programming; CIS 64C/G — Introduction to PL/SQL; CIS 44F — Introduction to Data Analytics; CIS 64E — Fundamentals of Large Scale Cloud Computing (AWS, GCP); CIS 102 — Ethical Hacking; CIS 104 — Digital Forensics; CIS 95F — Managing Cloud Projects.

## **ADDITIONAL INFORMATION**

---

- FAA Commercial Drone License; FAA Private Pilot License (intended CFI)
- O'Reilly & OER Publications: Data Analytics and Managing Cloud Projects (forthcoming Summer 2026)
- First Tee Assistant Golf Coach — San Francisco Chapter (2012–Present), Pittsburgh Chapter (2011–2012)
- Interests: Classical & Jazz Saxophone (Alto/Tenor/Baritone), Triathlon, Golf, Intramural Basketball & Soccer