

# O. FELIX AMORUWA

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

Technical Product Leader with 12+ years building agentic AI platforms, developer-facing SDKs, and enterprise-scale infrastructure — from shipping a multi-agent orchestration framework with MCP connectors and governed tool invocation (StreamIO) to scaling a self-service developer platform to 675M+ engagements at 50K TPS (Intuit). Built production AI evaluation frameworks with automated safety gates and CI/CD quality pipelines; NeurIPS published researcher with hands-on experience across RAG retrieval, RL post-training, and LLM-powered applications. Proven cross-functional leader who has driven platform strategy to CTO and VP level, reducing developer time-to-production from weeks to minutes.

## PROFESSIONAL EXPERIENCE

---

### STREAMIO AI | Oakland, CA

Founder & CEO — Agentic AI Platform (macOS, Linux, iOS & Web) 09/24 – Present

- Designed and shipped OpenClaw multi-agent orchestration framework with gateway protocol, subagent delegation, profile management, and session switching — providing a managed agent runtime enabling coordinated agentic workflows across real estate, insurance, health/dental, and financial markets domains.
- Implemented MCP SDK (Claude) integration exposing screen capture tools to AI coding assistants via a standardized MCP server — direct hands-on experience with Model Context Protocol connector architecture and governed tool invocation.
- Built cross-platform developer SDK with embedded terminal (xterm.js/node-pty, 20+ slash commands), native macOS ScreenCaptureKit Swift bridge, and shipped signed/notarized desktop builds (macOS DMG, Linux .deb) — delivering a self-service developer experience from prototype to production.
- Engineered end-to-end auth and payments pipeline (Kinde OAuth 2.0, Stripe tiered subscriptions, Electron SafeStorage) with ElevenLabs TTS/STT integration — enterprise-grade security, credential management, and reliability across all platform tiers.
- Led 0-to-1 product strategy, agentic AI development, and go-to-market execution including customer discovery and iterative product refinement.

### FINTELLECT AI | Oakland, CA

Founder & CEO — AI Financial Platform (iOS/Android/Web) 09/24 – Present

- Architected RAG retrieval pipeline with ChromaDB vector store, multi-provider LLM orchestration (Claude, GPT-4, Gemini) with model gateway abstraction and fallback routing, structured output validation, and token budget optimization — building the intelligence layer for a production agentic application.
- Built domain-specific Fintellect Agents — conversational AI agents scoped to distinct financial focal points, delivering context-aware advisory interactions with configurable guardrails and per-agent context management.
- Executed go-to-market through App Store launch following customer discovery and iterative refinement with trader feedback.

### INTUIT | Mountain View, CA

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

- Achieved 275% YoY growth in ICE platform engagements, scaling to 675M+ in FY23 across QuickBooks, TurboTax, Mint, Mailchimp, and Credit Karma; scaled throughput from 6K to 50K TPS via rSocket migration supporting ~1.5M concurrent connections with sub-25ms TP99 — enterprise-grade reliability at scale.
- Extended Java and Python SDK Starter Kits with scaffolding templates, build configurations (Gradle/Maven), testing frameworks, and CI/CD integration — empowering developers to go from zero to production-ready microservice in minutes; measured success by adoption and developer onboarding velocity.
- Delivered ICE Self-Service platform (DevPortal, GitOps config, ICE Playground), reducing developer time-to-production from 2–3 weeks to minutes in pre-prod and <24 hours for production — self-service by construction, mitigating \$1M+ in projected opex growth.
- Conducted enterprise-wide Service Language Assessment across 9 languages, analyzing usage data and developer feedback to inform strategic platform investment decisions presented to CTO; built Asterias, a declarative asset lifecycle management platform with GraphQL API.
- Initiated MSaaS Drift Detection and Resolution program: authored Java JAR library to scan Git repos for configuration drift, partnered with Design on DevPortal UI, and built remediation roadmap using OpenRewrite — gold-standard promotion pipeline from prototype to production-hardened service.
- Led Mailchimp GCP-to-AWS migration for MSaaS, delivering Golang service template, MySQL persistence integration, and updated DevPortal documentation — enterprise SaaS platform consolidation and migration.

### SPLUNK INC. | San Francisco, CA

Senior Product Manager — Search Orchestration 01/19 – 01/21

- Owned Search Service (Go microservices), Search Catalog (PostgreSQL metadata service), and Splunk Processing Language (SPL/SPL2) — built product roadmaps, PRDs, and acceptance criteria for Splunk Cloud Services; delivered Scheduler Service end-to-end in ~4 months, demoed at .conf19.
- Led query performance optimization initiative achieving up to 10x improvements in SCS search; designed RICE-based prioritization framework across 3 microservice backlogs balancing Fortune 500 customer, internal partner, and third-party developer requirements.

### 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 volume, 200+ internal enterprise customers) and ITSI Application Monitoring-as-a-Service; built Redis/XC10 caching capability for scalability, fault tolerance, and data redundancy across the enterprise.
- 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 software; 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 5 core eval types (factuality, reasoning, instruction-following, safety, code generation), adversarial safety testing with refusal detection, and CI/CD integration with regression detection and automated safety gates — no agent reaches production without passing quality thresholds.
- Statistical rigor: bootstrap confidence intervals, Welch's t-test, Cohen's d effect size, saturation detection. Stack: FastAPI orchestrator, TimescaleDB, Redis job queue, Next.js dashboard, Ollama — LLM-as-judge evaluation pipeline with domain-specific scoring.

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

- Built 3-phase RL post-training workbench: Reward Lab for A/B testing reward functions across 4 datasets, Playground for real TRL-powered GRPO/DPO training with live SSE metric streaming, and Arena for head-to-head framework benchmarking (TRL, VeRL, OpenRLHF, NeMo RL) with GPU Docker passthrough — standardized evaluation framework across AI training pipelines.
- Implemented 12 RL algorithms (PPO, GRPO, DAPO, DPO, SimPO, IPO, KTO, ORPO, and more) with algorithm-specific metric profiles and standardized throughput/memory/convergence benchmarking; MLflow-compatible experiment tracking architecture.

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

- Built automated visual evaluation pipeline reducing evaluation cycles from 72 hours to ~4 minutes; multimodal AI (Claude/GPT-4V) performs spatial reasoning on prediction frames, generating structured PASS/FAIL reports with confidence scores — LLM-as-judge scoring with zero-integration architecture via screen capture.

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

- Production ML platform in PyTorch with 5 neural architectures, MLflow experiment tracking, Optuna HPO, and FastAPI serving — 823 automated tests, Docker orchestration (6 containers), spanning 413 to 8B parameters. NeurIPS 2014 accepted paper on neural networks for protein secondary structure prediction.

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/64G — PL/SQL; CIS 44F — Data Analytics; CIS 64E — 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