

O. FELIX AMORUWA

famoruwa@berkeley.edu | 909-731-9011 | felixamoruwa.info

Technical Product Leader with 12+ years building agentic AI platforms, developer-facing SDKs, and enterprise infrastructure at scale — from shipping self-service developer platforms serving 675M+ engagements at Intuit to building multi-agent orchestration frameworks with MCP connectors and RAG intelligence layers as a founder. Deep hands-on experience with agent runtimes, LLM evaluation pipelines, model gateway abstraction, and governed tool invocation. Proven ability to drive cross-functional platform strategy from 0-to-1 through enterprise scale, communicate roadmaps at CTO/VP level, and reduce time-to-production from weeks to minutes. NeurIPS published researcher; MLflow practitioner.

PROFESSIONAL EXPERIENCE

STREAMIO AI | Oakland, CA

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

- Architected OpenClaw multi-agent orchestration framework with gateway protocol, subagent delegation, profile management, and session switching — a managed agent runtime enabling governed, multi-step orchestration across real estate, insurance, health/dental, and financial markets domains.
- Implemented MCP SDK (Claude Model Context Protocol) integration and built MCP server exposing screen capture tools to AI coding assistants — direct experience with standardized MCP connectors and governed tool invocation.
- Shipped self-service developer tooling: embedded terminal (xterm.js/node-pty, 20+ slash commands), cross-platform SDK with native macOS ScreenCaptureKit (Swift), and end-to-end auth/payments pipeline (Kinde OAuth 2.0, Stripe tiered plans, Electron SafeStorage) — reducing developer integration from days to minutes.
- Built production Electron + React + TypeScript desktop application with 100+ components, Redux Toolkit state management, and real-time AI analysis pipeline using Claude — enabling contextual conversations, Pine Script generation, and automated CMA reports via Redfin/Zillow APIs.
- Led 0-to-1 product strategy, agentic AI development, and go-to-market execution including customer discovery and iterative platform refinement.

FINTELLECT AI | Oakland, CA

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

- Architected RAG intelligence layer with ChromaDB vector store, multi-provider LLM model gateway (Claude, GPT-4, Gemini) with fallback routing, structured output validation, and token budget optimization — direct analog to Databricks' three-layer intelligence architecture.
- Built domain-specific Fintellect Agents — conversational AI agents scoped to distinct financial focal points, delivering context-aware advisory interactions with governed tool invocation and configurable per-agent scope.
- Integrated AI-powered charting and macroeconomic analysis tools with LLM models for automated trade analysis; executed go-to-market through App Store launch.

INTUIT | Mountain View, CA

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

- Achieved 275% YoY growth in ICE 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 platform reliability at 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 — self-service by construction, mitigating \$1M+ in projected opex growth.
- 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 velocity.
- 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: wrote Java JAR library to scan Git repos for configuration drift and built remediation roadmap using OpenRewrite — platform governance and gold-standard promotion pipeline analog.
- Led Mailchimp GCP-to-AWS migration for MSaaS; implemented ICE Presence in async chat generating \$480K/month in additional invoicing. DeveloperWeek 2022 speaker.

SPLUNK INC. | San Francisco, CA

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

- Owned Search Service (Go microservices), Search Catalog (PostgreSQL metadata service), and Splunk Processing Language (SPL/SPL2) — built PRDs, user stories, and acceptance criteria for Splunk Cloud Services; led query performance optimization achieving up to 10x improvements.
- Delivered Scheduler Service end-to-end in ~4 months, enabling scheduled search capabilities for first-party applications; designed RICE-based prioritization framework across 3 microservice backlogs balancing Fortune 500 and third-party developer requirements. Splunk .conf18/.conf19 speaker.

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); built Redis caching capability across the enterprise for scalability and fault tolerance; led demand forecasting and capacity planning 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 improving case resolution time 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 automated safety gates in CI/CD — a gold-standard promotion pipeline ensuring no agent reaches production without passing quality and safety thresholds.

- Statistical rigor: bootstrap confidence intervals, Welch's t-test, Cohen's d effect size, saturation detection, and regression detection. Stack: FastAPI orchestrator, TimescaleDB, Redis job queue, Next.js dashboard, Ollama — direct analog to Databricks' LLM-as-judge evaluation framework.

RL Workbench — Post-Training RL Platform (2026)

- Built 3-phase RL post-training workbench: Reward Lab (A/B testing reward functions across GSM8K, MATH, HumanEval, UltraFeedback), Playground (real TRL-powered GRPO/DPO training with live SSE metric streaming on MPS/CUDA), and Arena (head-to-head framework benchmarking: TRL, VeRL, OpenRLHF, NeMo RL) with GPU passthrough in Docker containers.
- Implemented 12 RL algorithms (PPO, GRPO, DAPO, DPO, SimPO, IPO, KTO, ORPO, SPPO, REINFORCE, REINFORCE++, RLOO) with standardized throughput/memory/convergence benchmarking — evaluation framework experience directly applicable to Databricks Mosaic AI model training and quality gates.

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

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

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

- Built zero-integration automated evaluation system using multimodal AI (Claude/GPT-4V) for spatial reasoning on prediction frames, generating structured PASS/FAIL reports with confidence scores — 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/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