

O. FELIX AMORUWA

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

Staff PM and hands-on AI builder with 12+ years driving technically complex platform products from inception to enterprise scale — including agentic AI systems, multi-agent orchestration frameworks, and AI/ML platform infrastructure. Scaled Intuit's ICE platform to 675M+ engagements (275% YoY) and 50K TPS; built OpenClaw multi-agent orchestration and a full RL post-training workbench benchmarking GRPO/DPO across TRL, VeRL, OpenRLHF, and NeMo RL. Trusted cross-functional AI expert who has presented platform strategy to CTO-level leadership and evangelized adoption across 30+ product SKUs. NeurIPS published researcher; hands-on in PyTorch, FastAPI, and SQL.

PROFESSIONAL EXPERIENCE

STREAMIO AI | Oakland, CA

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

- Architected and shipped OpenClaw multi-agent orchestration framework with gateway protocol, subagent delegation, profile management, and session switching — enabling coordinated agentic AI workflows across real estate, insurance, health/dental, and financial markets verticals; directly informs context engineering and multi-agent platform design.
- Built production Electron + React + TypeScript desktop application with 100+ components and Redux Toolkit state management, integrating Claude via MCP SDK for real-time contextual AI analysis — delivering agentic decision-making capabilities to end users across multiple industry verticals.
- Shipped MCP server exposing screen capture tools to AI coding assistants, embedded terminal (xterm.js/node-pty, 20+ slash commands), and cross-platform builds (macOS DMG with code signing/notarization, Linux .deb) — demonstrating end-to-end 0-to-1 AI product execution.
- Led 0-to-1 product strategy, roadmap definition, customer discovery, and go-to-market execution — iterating on platform offerings based on user feedback and identifying new use-case opportunities to drive adoption.

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 fallback routing, structured output validation, and token budget optimization — accelerating data-driven decision-making for retail investors.
- Built domain-specific conversational AI agents embedded in mobile and web app, each scoped to distinct financial focal points — delivering context-aware advisory interactions and demonstrating agentic platform design at the application layer.
- Executed go-to-market through App Store launch, led customer discovery, and established industry partnerships — evangelizing platform adoption and refining product-market fit based on trader feedback.

INTUIT | Mountain View, CA

Staff Product Manager — AI/ML Platform Infrastructure & Developer Frameworks 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.
- 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 — mitigating \$1M+ in projected opex growth and driving measurable platform adoption across 30+ product SKUs.
- Leveraged SQL and BigQuery telemetry to track KPIs, prioritize developer pain points, and build data-driven roadmaps across ~20 mobile apps; presented enterprise-wide Service Language Assessment (9 languages) and strategic investment recommendations directly to CTO.
- Implemented ICE Presence in async chat, generating \$480K/month in additional invoicing; deployed Background-to-Foreground Messaging on iOS/Android with <100ms latency — showcasing AI/ML platform capabilities to leadership at DeveloperWeek 2022.
- Initiated MSaaS Drift Detection program: authored Java JAR library to scan Git repos for configuration drift and built remediation roadmap using OpenRewrite — grounding build-vs-buy decisions in data and engineering research.
- Led Mailchimp GCP-to-AWS migration for MSaaS, delivering Golang service template, MySQL persistence integration, and updated DevPortal documentation to meet production deadline.

SPLUNK INC. | San Francisco, CA

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

- Owned Search Service (Go microservices), Search Catalog (PostgreSQL metadata service), and Splunk Processing Language (SPL/SPL2) — defined roadmaps, PRDs, and acceptance criteria for Splunk Cloud Services, managing engineers and scientists toward complex technical milestones.
- Delivered Scheduler Service end-to-end in ~4 months (June–Oct 2019), enabling scheduled search capabilities for first-party applications; demoed at Splunk .conf19 to 10,000+ attendees.
- Led query performance optimization for Fortune 500 beta customer (Assurance), achieving up to 10x improvements in SCS search — grounding technology bets in benchmark data and customer requirements.
- Designed repeatable RICE-based prioritization framework across 3 microservice backlogs, balancing internal partner, third-party developer, and Fortune 500 customer requirements; presented at Splunk .conf18.

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 — driving internal platform adoption at scale.
- Built Redis and XC10 caching capability enterprise-wide to address scalability, fault tolerance, and data redundancy; 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 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 Monte Carlo simulation model to optimize DMAIC phase distribution across a \$494M portfolio in Global Technology & Operations — applying quantitative analysis to data-driven decision-making.

AI/ML RESEARCH & PROJECTS

RL Workbench — Post-Training RL Platform (2026)

- Built 3-phase RL post-training workbench covering the full RLHF/DPO pipeline: Reward Lab for A/B testing reward functions (RLVR, learned, hybrid) across 4 datasets, Playground for real TRL-powered GRPO/DPO training with live SSE metric streaming on Apple Silicon (MPS) or CUDA, and Arena for head-to-head framework benchmarking (TRL, VerL, OpenRLHF, NeMo RL) with GPU passthrough in Docker containers.
- Implemented 12 RL algorithms (PPO, GRPO, DAPO, REINFORCE, REINFORCE++, RLOO, DPO, SimPO, IPO, KTO, ORPO, SPPO) with algorithm-specific metric profiles and standardized throughput/memory/convergence benchmarking — directly applicable to evaluating AI/ML platform technology bets.

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 automated safety gates and regression detection.
- 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 — enabling automated KPI measurement of AI model quality.

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

- Built automated visual evaluation system repurposing multimodal AI pipeline (Claude/GPT-4V) to score model outputs against natural-language rubrics — reducing evaluation cycles from 72 hours to ~4 minutes with zero-integration screen-capture architecture.

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

- Production ML platform in PyTorch with 5 neural architectures (feedforward, GRU, Transformer, ESM-2, multi-task), MLflow experiment tracking, Optuna HPO, and FastAPI serving — 823 automated tests, Docker orchestration (6 containers); NeurIPS 2014 accepted paper.

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