

O. FELIX AMORUWA

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

Forward Deployed Product Manager and AI Agent builder with 12+ years owning complex deployments end-to-end — from pre-sale scoping through post-go-live optimization. Designed and shipped multi-agent orchestration frameworks (OpenClaw) and domain-specific conversational AI agents across real estate, financial, and insurance verticals, with hands-on LLM orchestration across Claude, GPT-4, and Gemini. Scaled AI platform deployments to 675M+ engagements at Intuit, driving measurable customer ROI. NeurIPS published researcher; UC Berkeley engineering graduate.

PROFESSIONAL EXPERIENCE

STREAMIO AI | Oakland, CA

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

- Designed and deployed OpenClaw multi-agent orchestration framework with gateway protocol, subagent delegation, profile management, and session switching — enabling coordinated AI agent workflows across real estate, insurance, health/dental, and financial markets verticals, directly analogous to enterprise AI Agent deployment lifecycle ownership.
- Led 0-to-1 AI Agent product strategy and go-to-market execution: conducted customer discovery interviews, defined agent use cases and success criteria, iterated on workflows based on user feedback, and drove deployment from concept to production.
- Built production AI agent application integrating Claude (MCP SDK) for real-time contextual analysis, Pine Script generation, and automated real estate CMA reports via Redfin/Zillow APIs — delivering measurable automation value across multiple customer segments.
- Shipped cross-platform production builds (macOS, Linux, iOS) with native integrations, embedded terminal with 20+ slash commands, and MCP server exposing screen capture tools to AI coding assistants — demonstrating hands-on, bias-for-action build style.
- Engineered real-time HLS livestreaming pipeline with multi-stream canvas compositing (up to 9 concurrent sources at 30fps), FFmpeg transcoding, and WebSocket communication layer — end-to-end technical ownership across the full stack.

FINTELLECT AI | Oakland, CA

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

- Built and deployed domain-specific conversational AI agents embedded in mobile and web app, each scoped to distinct financial focal points — delivering guided, context-aware advisory interactions that drive user automation and engagement.
- Architected multi-provider LLM orchestration pipeline (Claude, GPT-4, Gemini) with RAG retrieval (ChromaDB), fallback routing, structured output validation, and token budget optimization — hands-on LLM-first product development.
- Led customer discovery and go-to-market execution including App Store launch; refined agent workflows based on trader feedback and established partnerships with industry influencers.

INTUIT | Mountain View, CA

Staff Product Manager — *Platform Infrastructure & Developer Experience* 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 — enterprise-scale deployment ownership with measurable customer ROI.
- 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.
- Conducted enterprise-wide Service Language Assessment across 9 languages, analyzing usage data and developer feedback to inform strategic decisions presented to CTO — senior stakeholder engagement at scale.
- Worked cross-functionally across ~20 mobile apps and 30+ product SKUs using telemetry and SQL/BigQuery to prioritize platform investments; scaled throughput from 6K to 50K TPS via rSocket migration.

SPLUNK INC. | San Francisco, CA

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

- Delivered Scheduler Service end-to-end in ~4 months, enabling scheduled search capabilities for first-party applications and demoed at Splunk .conf19 — fast, customer-facing deployment execution.
- Led query performance optimization for beta enterprise customer (Assurance), achieving up to 10x performance improvements — direct customer ROI ownership with cross-functional engineering partnership.
- Designed RICE-based prioritization framework balancing Fortune 500 customer requirements across 3 microservice backlogs; owned PRDs, user stories, and acceptance criteria for Splunk Cloud Services.

KAISER PERMANENTE | Pleasanton, CA

SOA Technical Product Manager 08/12 – 12/18

- Led enterprise rollout of Splunk Logging-as-a-Service (1.7 TB daily volume, 200+ internal enterprise customers) and ITSI Application Monitoring-as-a-Service — large-scale platform deployment and customer success ownership.

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 products; led cross-functional root cause analysis and customer-facing resolution, improving case resolution time by 20%.

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.

AI/ML RESEARCH & PROJECTS

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

- Repurposed StreamIO's multi-agent AI pipeline to score model outputs against natural-language evaluation rubrics using multimodal AI (Claude/GPT-4V) — reducing evaluation cycles from 72 hours to ~4 minutes; zero-integration architecture captures from any visualization tool via screen capture.

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 regression gates. Stack: FastAPI, TimescaleDB, Redis, Next.js, Ollama.

RL Workbench — Post-Training RL Platform (2026)

- Built 3-phase RL post-training workbench benchmarking GRPO/DPO across TRL, VeRL, OpenRLHF, and NeMo RL; implemented 12 RL algorithms (PPO, GRPO, DAPO, DPO, SimPO, and more) with live SSE metric streaming on Apple Silicon (MPS) or CUDA.

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

- NeurIPS 2014 accepted paper; production PyTorch platform with 5 neural architectures, MLflow, Optuna HPO, and FastAPI serving — spanning 413 to 8B parameters (19M-fold scale increase from original hand-coded C++ BPTT system).

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 — 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

- Java Programming, PL/SQL, Data Analytics, Large Scale Cloud Computing (AWS, GCP), Ethical Hacking, Digital Forensics, 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)
- Interests: Classical & Jazz Saxophone (Alto/Tenor/Baritone), Triathlon, Golf, Intramural Basketball & Soccer