

O. FELIX AMORUWA

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

Senior Product Manager with 12+ years leading customer-facing platform products and AI chat solutions at scale — including owning Intuit's chat platform to 675M+ engagements and shipping self-service tooling that cut onboarding from weeks to minutes. Built production conversational AI agents with multi-LLM orchestration (Claude, GPT-4, Gemini), RAG pipelines, and multi-agent frameworks for customer-facing workflows. Proven track record delivering end-to-end product features, managing third-party vendor integrations, and using data to optimize CX metrics across globally distributed teams. NeurIPS published; BS Computational Engineering, UC Berkeley.

PROFESSIONAL EXPERIENCE

INTUIT | Mountain View, CA

Staff Product Manager — Chat Platform, Self-Service & Developer Infrastructure 06/21 – 09/24

- Owned ICE chat platform and self-service tooling, achieving 275% YoY engagement growth and scaling to 675M+ engagements in FY23 across QuickBooks, TurboTax, Mint, Mailchimp, and Credit Karma — driving measurable CSAT and deflection improvements at consumer scale.
- Implemented ICE Presence in async chat, generating \$480K/month in additional invoicing; deployed Background-to-Foreground Messaging on iOS/Android with <100ms latency, improving real-time customer communication across 20+ mobile apps.
- Delivered ICE Self-Service platform (DevPortal, GitOps config, ICE Playground), reducing onboarding from 2–3 weeks to minutes in pre-prod and <24 hours for production — mitigating \$1M+ in projected opex growth through automation.
- Led Mailchimp GCP-to-AWS platform migration for MSaaS, coordinating cross-functional teams across Engineering, Program Management, and third-party vendors to meet production deadline.
- Used SQL and BigQuery telemetry to analyze usage data across 30+ product SKUs, prioritizing roadmap investments and surfacing customer pain points to inform data-driven product decisions.
- Collaborated with globally distributed teams across 9 engineering language ecosystems; presented strategic recommendations directly to CTO based on enterprise-wide usage analysis.

FINTELLECT AI | Oakland, CA

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

- Built and launched conversational AI agents scoped to distinct customer focal points — delivering guided, context-aware self-service advisory interactions across mobile and web, with multi-provider LLM orchestration (Claude, GPT-4, Gemini) and fallback routing.
- Architected RAG retrieval pipeline with ChromaDB vector store, structured output validation, and token budget optimization — enabling accurate, low-latency GenAI responses for customer-facing chat.
- Led end-to-end product lifecycle: customer discovery interviews, iterative feature refinement based on user feedback, and App Store launch — applying CX metrics to guide roadmap prioritization.

STREAMIO AI | Oakland, CA

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

- Implemented OpenClaw multi-agent orchestration framework with gateway protocol and subagent delegation — enabling coordinated AI agent workflows for customer-facing use cases across real estate, insurance, health/dental, and financial markets.
- Integrated Claude (MCP SDK) for real-time contextual AI conversations and automated report generation via third-party APIs (Redfin/Zillow), demonstrating end-to-end GenAI feature delivery with external vendor integrations.
- Built end-to-end auth and payments pipeline using Kinde OAuth 2.0, Stripe API integrations (tiered subscriptions), and secure credential management — shipping production-ready features from 0 to 1.

SPLUNK INC. | San Francisco, CA

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

- Authored detailed PRDs, user stories, and acceptance criteria for Splunk Cloud Services (SCS) microservices — delivering clear, actionable requirements for both technical and non-technical stakeholders.
- Delivered Scheduler Service end-to-end in ~4 months (June–Oct 2019), enabling scheduled capabilities for first-party applications; demoed at Splunk .conf19.
- Designed RICE-based prioritization framework across 3 microservice backlogs, balancing Fortune 500 customer requirements, internal partner needs, and third-party developer integrations.

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 customers) and ITSI Application Monitoring-as-a-Service — managing third-party platform integrations and cross-departmental adoption.
- Led demand forecasting and capacity planning for SOA product suite across multiple datacenters, using data to inform infrastructure investment decisions.

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 software products, improving case resolution time by 20% — demonstrating deep focus on customer satisfaction and CX metric improvement.

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

aeval — AI Model Evaluation Platform (2025–2026)

- Built local-first AI 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. Stack: FastAPI orchestrator, TimescaleDB, Redis job queue, Next.js dashboard, 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 with standardized throughput/memory/convergence benchmarking and GPU passthrough in Docker containers.

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

- NeurIPS 2014 accepted paper. Original 2004 system: hand-coded neural network in C++ with custom BPTT; 2026 rewrite in PyTorch spans 413 to 8B parameters across 5 architectures with MLflow, Optuna HPO, and FastAPI serving.

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