

O. FELIX AMORUWA

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

Technical Product Leader with 12+ years building observability platforms, telemetry-driven developer experiences, and AI-powered insights at scale — from owning Splunk's Search Orchestration and SPL query language to scaling Intuit's ICE platform to 675M+ engagements with sub-25ms TP99 latency. Proven track record translating raw telemetry into actionable, proactively surfaced insights that drive customer action on performance, reliability, and cost. NeurIPS published researcher; firsthand experience operating GPU-powered AI workloads as a builder of RL post-training benchmarking infrastructure.

PROFESSIONAL EXPERIENCE

SPLUNK INC. | San Francisco, CA

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

- Owned Search Service (Go microservices), Search Catalog (PostgreSQL metadata service), and Splunk Processing Language (SPL/SPL2) — the core query and telemetry pipeline powering Splunk Cloud Services observability experiences for Fortune 500 customers.
- Delivered Scheduler Service end-to-end in ~4 months, enabling scheduled search and automated alerting capabilities for first-party applications; demoed at Splunk .conf19.
- Led query performance optimization initiative for beta customer (Assurance), building mirrored Enterprise topology for benchmark testing and achieving up to 10x performance improvements — directly improving workload efficiency and reliability signal quality.
- Designed repeatable RICE-based prioritization framework across 3 microservice backlogs, balancing internal partner, third-party developer, and Fortune 500 customer requirements. Splunk .conf18 Speaker.

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 — deep infrastructure performance and reliability at hyperscale.
- Leveraged SQL and BigQuery telemetry to surface high-value signals and prioritize developer pain points across ~20 mobile apps and 30+ product SKUs; built Asterias, a declarative asset lifecycle management platform with GraphQL API.
- Initiated MSaaS Drift Detection and Resolution program — wrote Java JAR library to proactively scan Git repos for configuration drift, partnered with Design on DevPortal UI, and built remediation roadmap using OpenRewrite; exemplifying proactive, automated signal surfacing over manual discovery.
- 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.
- Conducted enterprise-wide Service Language Assessment across 9 languages, analyzing usage data and developer feedback to inform strategic investment decisions presented 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. DeveloperWeek 2022 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) and ITSI Application Monitoring-as-a-Service — end-to-end ownership of large-scale observability platform products.
- Built caching capability using Redis and XC10 across the enterprise to address scalability, fault tolerance, and data redundancy challenges.
- Led demand forecasting and capacity planning for SOA product suite, analyzing demand drivers and developing IT capacity plans across multiple datacenters.

STREAMIO AI | Oakland, CA

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

- Implemented OpenClaw multi-agent orchestration framework with gateway protocol and subagent delegation — enabling AI-powered automated analysis and natural-language insight delivery across multiple industry verticals.
- Built real-time AI analysis pipeline using Claude (MCP SDK) for contextual signal interpretation from screen captures; engineered HLS livestreaming pipeline with multi-stream canvas compositing at 30fps and WebSocket communication layer.
- Led 0-to-1 product strategy, AI development, and go-to-market execution including customer discovery interviews and iterative product refinement.

FINTELLECT AI | Oakland, CA

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

- Architected RAG retrieval pipeline with ChromaDB vector store and multi-provider LLM orchestration (Claude, GPT-4, Gemini) with fallback routing — translating complex financial telemetry into actionable, AI-powered insights for retail investors.
- Built AI-powered charting and macroeconomic analysis tools surfacing cost optimization and risk signals; integrated LLM models for automated trade analysis at scale.

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 products; 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

RL Workbench — Post-Training RL Platform (2026)

- Built GPU-powered post-training workbench with live SSE metric streaming, standardized throughput/memory/convergence benchmarking, and cross-tab workflow lineage tracking across TRL, VeRL, OpenRLHF, and NeMo RL — firsthand experience operating and observing the AI workloads CoreWeave customers run at scale.
- Implemented 12 RL algorithms (PPO, GRPO, DAPO, REINFORCE, DPO, SimPO, and more) with GPU Docker passthrough, algorithm-specific metric profiles, and head-to-head framework benchmarking — producing actionable performance and cost signals across training runs.

aeval — AI Model Evaluation Platform (2025–2026)

- Built local-first evaluation platform with automated safety gates, CI/CD regression detection, and statistical rigor (bootstrap confidence intervals, Welch's t-test, Cohen's d) — proactively surfacing high-signal, low-noise quality metrics across 5 eval types.
- Stack: FastAPI orchestrator, TimescaleDB, Redis job queue, Next.js dashboard, Ollama — production observability architecture for AI model performance and reliability.

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); scales from 413 to 8B parameters.
- NeurIPS 2014 accepted paper on artificial neural networks for protein secondary 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 — reducing evaluation cycles from 72 hours to ~4 minutes with structured PASS/FAIL reports and confidence scores.

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