

O. FELIX AMORUWA

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

Staff PM with 12+ years translating raw telemetry into actionable, customer-facing observability experiences — from owning Splunk's Search Orchestration and SPL/SPL2 query platform to scaling ICE platform telemetry to 675M+ engagements at 50K TPS with sub-25ms TP99 at Intuit. Built AI-powered insights pipelines, multi-LLM orchestration with automated signal discovery, and a GPU workbench benchmarking RL frameworks across CUDA clusters — giving me firsthand fluency in the workload efficiency and cost optimization signals CoreWeave's customers need to act on. NeurIPS published; UC Berkeley BS Computational Engineering Science.

PROFESSIONAL EXPERIENCE

SPLUNK INC. | San Francisco, CA

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

- Owned Search Service (Go microservices), Search Catalog (PostgreSQL metadata service), and Splunk Processing Language (SPL/SPL2) — defining roadmaps, PRDs, and acceptance criteria for Splunk Cloud Services; directly analogous to PromQL/LogQL-based observability query pipelines.
- Led query performance optimization initiative for enterprise beta customer (Assurance), building mirrored topology for benchmark testing and achieving up to 10x performance improvements — translating raw metrics and logs into high-signal, low-noise customer experiences.
- Delivered Scheduler Service end-to-end in ~4 months (June–Oct 2019), enabling scheduled search and alerting capabilities for first-party applications; demoed at Splunk .conf19.
- Designed repeatable RICE-based prioritization framework across 3 microservice backlogs, balancing Fortune 500 customer requirements with internal platform needs. Splunk .conf18 Speaker.

KAISER PERMANENTE | Pleasanton, CA

SOA Technical Product Manager — Observability & Platform Infrastructure 08/12 – 12/18

- Led development and enterprise rollout of Splunk Logging-as-a-Service (1.7 TB daily log volume, 200+ internal enterprise customers) and ITSI Application Monitoring-as-a-Service — owning the full observability platform lifecycle from telemetry ingestion to customer-facing dashboards and alerts.
- Built caching capability using Redis and XC10 across the enterprise to address scalability, fault tolerance, and data redundancy — improving reliability signals for downstream monitoring consumers.
- Led demand forecasting and capacity planning for SOA product suite, analyzing demand drivers and developing IT capacity plans across multiple datacenters — translating infrastructure metrics into cost and workload optimization decisions.

INTUIT | Mountain View, CA

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

- Scaled ICE platform to 675M+ engagements in FY23 across QuickBooks, TurboTax, Mint, Mailchimp, and Credit Karma; drove 50K TPS throughput via rSocket migration supporting ~1.5M concurrent connections with sub-25ms TP99 — owning performance and reliability metrics at hyperscale.
- Used SQL and BigQuery telemetry to surface actionable insights across ~20 mobile apps and 30+ product SKUs, directly informing roadmap prioritization and developer experience investments.
- Built Asterias, a declarative asset lifecycle management platform with GraphQL API — translating raw configuration and usage data into proactively surfaced signals for platform health and drift detection.
- Initiated MSaaS Drift Detection and Resolution program: wrote Java JAR library to scan Git repos for configuration drift and built remediation roadmap — proactive observability for infrastructure-as-code 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, 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. DeveloperWeek 2022 speaker.

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 generation across multiple industry verticals.
- Built production Electron + React + TypeScript desktop application with real-time AI analysis of screen captures using Claude (MCP SDK), enabling contextual, proactively surfaced insights from live data streams.
- Engineered real-time HLS livestreaming pipeline with multi-stream canvas compositing (up to 9 concurrent sources at 30fps) and WebSocket communication layer — high-throughput telemetry pipeline experience directly applicable to GPU workload observability.

FINTELLECT AI | Oakland, CA

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

- Architected RAG retrieval pipeline with multi-provider LLM orchestration (Claude, GPT-4, Gemini), fallback routing, structured output validation, and token budget optimization — AI-powered insight surfacing with automated signal discovery.
- Built AI-powered charting and macroeconomic analysis tools for cost optimization, asset class selection, and risk management — translating complex financial telemetry into actionable, customer-facing insights.

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 improved 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-cluster post-training workbench benchmarking GRPO/DPO across TRL, VeRL, OpenRLHF, and NeMo RL with live SSE metric streaming on MPS/CUDA and Docker GPU passthrough — firsthand experience operating the exact GPU-powered AI workloads CoreWeave's customers run, with real-time telemetry, throughput/memory/convergence benchmarking, and cross-tab workflow lineage tracking.

- Implemented 12 RL algorithms (PPO, GRPO, DAPO, REINFORCE, REINFORCE++, RLOO, DPO, SimPO, IPO, KTO, ORPO, SPPO) with algorithm-specific metric profiles and standardized benchmarking across frameworks — translating raw training telemetry into actionable performance and cost signals.

aeval — AI Model Evaluation Platform (2025–2026)

- Built local-first evaluation platform with 5 eval types, adversarial safety testing with automated refusal detection, and CI/CD regression gates — proactive, high-signal alerting with statistical rigor: bootstrap confidence intervals, Welch's t-test, Cohen's d effect size, saturation detection.
- Stack: FastAPI orchestrator, TimescaleDB (time-series metrics), Redis job queue, Next.js dashboard, Ollama — directly applicable to observability infrastructure for AI workloads.

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

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

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); original 2004 system hand-coded in C++ with custom BPTT.

NeurIPS 2014 — Montreal, Canada

- Accepted paper on artificial neural networks for protein secondary structure prediction.

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)
- Interests: Classical & Jazz Saxophone (Alto/Tenor/Baritone), Triathlon, Golf, Intramural Basketball & Soccer