

O. FELIX AMORUWA

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

Technical Product Manager with 12+ years owning distributed data platforms, developer tooling, and ML infrastructure at scale — from scaling Intuit's platform to 675M+ engagements and 50K TPS to hand-building a distributed RL post-training workbench benchmarking TRL, VeRL, OpenRLHF, and NeMo RL today. Deep experience driving developer experience, open source ecosystem growth, and enterprise customer engagement across data processing and AI infrastructure products. NeurIPS published ML researcher. Bay Area based.

PROFESSIONAL EXPERIENCE

INTUIT | Mountain View, CA

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

- Achieved 275% YoY growth in 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 — demonstrating ownership of distributed data infrastructure at enterprise scale.
- Extended Java and Python SDK Starter Kits with scaffolding templates, build configurations (Gradle/Maven), testing frameworks, and CI/CD integration — driving developer experience improvements that took developers from zero to production-ready microservice in minutes.
- 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 (Java, Python, Kotlin, Go, TypeScript, Scala, PHP, C++, Groovy), analyzing usage data and developer feedback to inform strategic language investment decisions presented to CTO.
- Worked closely with telemetry and usage data (SQL, BigQuery) to prioritize developer pain points across ~20 mobile apps and 30+ product SKUs; built Asterias, a declarative asset lifecycle management platform with GraphQL API.
- Led Mailchimp GCP-to-AWS migration for MSaaS, delivering Golang template, MySQL persistence integration, and updated DevPortal documentation; initiated MSaaS Drift Detection program using Java JAR library to scan Git repos for configuration drift.
- 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.

SPLUNK INC. | San Francisco, CA

Senior Product Manager — Search Orchestration & Distributed Data Processing 01/19 – 01/21

- Owned Search Service (Go microservices), Search Catalog (PostgreSQL metadata service), and Splunk Processing Language (SPL/SPL2) — driving product roadmaps, PRDs, and acceptance criteria for scalable data processing capabilities in Splunk Cloud Services.
- Led query performance optimization initiative for enterprise beta customer, building mirrored topology for benchmark testing and achieving up to 10x performance improvements in distributed search — directly analogous to Ray Data's batch processing optimization challenges.
- Delivered Scheduler Service end-to-end in ~4 months, enabling scheduled data processing for first-party applications; demoed at Splunk .conf19.
- Designed repeatable RICE-based prioritization framework for 3 microservice backlogs, balancing internal developer, third-party ecosystem, and Fortune 500 enterprise customer requirements. Splunk .conf18 Speaker.

STREAMIO AI | Oakland, CA

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

- Led 0-to-1 product strategy, AI development, and go-to-market execution for a production AI platform — including customer discovery, iterative product refinement, and commercial launch.
- Implemented OpenClaw multi-agent orchestration framework with gateway protocol and subagent delegation — distributed AI agent workflows across multiple industry verticals.
- Engineered real-time HLS livestreaming pipeline with multi-stream canvas compositing (up to 9 concurrent sources at 30fps), FFmpeg transcoding, and WebSocket communication layer — scalable data processing architecture for real-time media workloads.

FINTELLECT AI | Oakland, CA

Founder & CEO — fintellectai.io (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, and token budget optimization — ML data infrastructure serving retail investors at scale.
- Built domain-specific conversational AI agents embedded in mobile and web app, delivering context-aware advisory interactions; executed go-to-market through App Store launch.

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) — large-scale data processing infrastructure serving a Fortune 500 healthcare organization.
- Built Redis caching capability across the enterprise to address scalability, fault tolerance, and data redundancy; led demand forecasting and capacity planning for SOA product suite 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 data products; led cross-functional root cause analysis improving case resolution time by 20% across the Support Organization.

AI/ML RESEARCH & PROJECTS

RL Workbench — Distributed ML Post-Training Platform (2026)

- Built 3-phase distributed RL post-training workbench covering the full RLHF/DPO pipeline: Reward Lab for designing and A/B testing reward functions across 4 datasets (GSM8K, MATH, HumanEval, UltraFeedback), 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 standardized throughput/memory/convergence benchmarking across frameworks — directly analogous to Ray Data's batch inference and ML training preprocessing workloads.

aeval — AI Model Evaluation Platform (2025–2026)

- Built local-first ML evaluation platform with 5 core eval types, adversarial safety testing, and CI/CD integration with automated safety gates. Stack: FastAPI orchestrator, TimescaleDB, Redis job queue, Next.js dashboard, Ollama — distributed data pipeline architecture for ML tooling.

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

- Built 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), 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 automated visual evaluation system reducing evaluation cycles from 72 hours to ~4 minutes; zero-integration architecture captures from any visualization tool via screen capture with multimodal AI scoring generating structured PASS/FAIL reports.

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

- Java Programming (CIS 35A/B, 36A/B), Introduction to PL/SQL (CIS 64C/G), Introduction to Data Analytics (CIS 44F), Fundamentals of Large Scale Cloud Computing — AWS, GCP (CIS 64E), Ethical Hacking (CIS 102), Digital Forensics (CIS 104), Managing Cloud Projects (CIS 95F).

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