

O. FELIX AMORUWA

famoruwa@berkeley.edu | 909-731-9011 | [felimamoruwa.info](mailto:famoruwa@berkeley.edu)

Technical PM and NeurIPS-published researcher with 12+ years bridging AI research and product delivery — uniquely positioned to translate safety research findings into guardrails and product-level controls. Built aeval, a production AI safety evaluation platform with adversarial testing, refusal detection, automated safety gates, and statistical regression detection. Hands-on experience with agentic AI systems, RAG pipelines, multi-agent orchestration, and RL post-training benchmarking across TRL, VeRL, OpenRLHF, and NeMo RL. Scaled enterprise AI platforms to 675M+ engagements at Intuit, building cross-functional alignment processes and proactive drift/regression detection programs.

AI SAFETY RESEARCH & PROJECTS

aeval — AI Model Evaluation Platform (2025–2026)

- Built local-first AI safety evaluation platform with 5 core eval types (factuality, reasoning, instruction-following, safety, code generation), adversarial safety testing with refusal detection, and data contamination detection via SHA-256 hashing — directly addressing the evaluation frameworks needed to assess how safety properties hold up as models evolve.
- Implemented statistical rigor throughout: bootstrap confidence intervals, Welch's t-test, Cohen's d effect size, and saturation detection; CI/CD integration with automated regression detection and safety gates to surface regressions before they reach production. Stack: FastAPI orchestrator, TimescaleDB, Redis job queue, Next.js dashboard, Ollama.

RL Workbench — Post-Training RL & Safety Alignment Platform (2026)

- Built 3-phase RL post-training workbench covering the full RLHF/DPO pipeline: Reward Lab for designing and A/B testing reward functions (RLVR, learned, hybrid) across 4 datasets (GSM8K, MATH, HumanEval, UltraFeedback) — directly analogous to safety reward modeling and alignment evaluation.
- Implemented 12 RL algorithms (PPO, GRPO, DAPO, DPO, SimPO, KTO, ORPO, and more) with cross-framework benchmarking (TRL, VeRL, OpenRLHF, NeMo RL) and standardized throughput/memory/convergence metrics — enabling the kind of behavioral research comparisons that inform safety roadmap decisions.

OpenClaw / StreamIO — Multi-Agent Orchestration & Agentic Safety (2024–2026)

- Implemented OpenClaw multi-agent orchestration framework with gateway protocol, subagent delegation, profile management, and session switching — gaining firsthand exposure to the unique safety challenges introduced by tool use, multi-step reasoning, and autonomous execution in agentic AI systems.
- Integrated Claude MCP SDK for real-time AI analysis and contextual agent workflows across regulated domains (real estate, financial markets, insurance) — surfacing practical misuse patterns and edge cases in production agentic deployments.

Fintellect AI — RAG Pipeline & LLM Safety in Regulated Domains (2024–2026)

- Architected RAG retrieval pipeline with ChromaDB vector store, multi-provider LLM orchestration (Claude, GPT-4, Gemini) with fallback routing, structured output validation, and token budget optimization — building direct familiarity with RAG poisoning vectors and LLM behavioral failure modes in high-stakes financial contexts.

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

- NeurIPS 2014 accepted paper on artificial neural networks for protein secondary structure prediction; 2026 rewrite spans 413 to 8B parameters across 5 architectures (feedforward, GRU, Transformer, ESM-2, multi-task) with MLflow experiment tracking, Optuna HPO, and FastAPI serving.
- Original 2004 system: hand-coded neural network in C++ with custom BPTT — establishing deep ML foundations that enable credible engagement with safety researchers on model behavior and evaluation methodology.

PROFESSIONAL EXPERIENCE

INTUIT | Mountain View, CA

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

- Achieved 275% YoY growth in ICE engagements, scaling to 675M+ in FY23 across QuickBooks, TurboTax, Mint, Mailchimp, and Credit Karma — demonstrating enterprise-scale PM execution and cross-functional alignment across research, engineering, and product teams.
- Initiated MSaaS Drift Detection and Resolution program: wrote Java JAR library to scan Git repos for configuration drift and built remediation roadmap — a proactive regression-detection process directly analogous to surfacing safety regressions before they reach customers.
- 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.
- Delivered ICE Self-Service platform (DevPortal, GitOps config, ICE Playground), reducing developer onboarding from 2–3 weeks to minutes; scaled throughput from 6K to 50K TPS via rSocket migration supporting ~1.5M concurrent connections with sub-25ms TP99.
- Conducted enterprise-wide Service Language Assessment across 9 languages, analyzing usage data and developer feedback to inform strategic decisions presented to CTO — building the written communication and executive translation skills essential for escalating safety findings.

SPLUNK INC. | San Francisco, CA

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

- Owned Search Service (Go microservices), Search Catalog (PostgreSQL metadata service), and Splunk Processing Language (SPL/SPL2) — built product roadmaps, PRDs, and acceptance criteria for Splunk Cloud Services, balancing internal research, third-party developer, and Fortune 500 customer requirements.
- Designed repeatable RICE-based prioritization framework for 3 microservice backlogs; led query performance optimization achieving up to 10x improvements for enterprise beta customers. Splunk .conf18 and .conf19 speaker.
- Delivered Scheduler Service end-to-end in ~4 months, enabling scheduled search capabilities for first-party applications — demonstrating ability to scope, ship, and demo research-adjacent product features under ambiguity.

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 in a regulated healthcare environment — building operational rigor in high-stakes, compliance-sensitive contexts.
- 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 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 software; led cross-functional root cause analysis and customer-facing resolution, improving case resolution time by 20% — establishing engineering foundations and escalation judgment that underpin credible engagement with safety researchers.

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 — quantitative risk modeling background applicable to safety risk prioritization frameworks.

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 and Hacker Investigation; CIS 95F — Managing Cloud Projects.

ADDITIONAL INFORMATION

- NeurIPS 2014 published researcher — artificial neural networks for protein secondary structure prediction; Lawrence Berkeley National Laboratory computational biology research under Dr. Steven Holbrook.
- O'Reilly & OER Publications: Data Analytics and Managing Cloud Projects (forthcoming Summer 2026); DeveloperWeek 2022 speaker.
- FAA Commercial Drone License; FAA Private Pilot License (intended CFI).
- First Tee Assistant Golf Coach — San Francisco Chapter (2012–Present); Interests: Classical & Jazz Saxophone (Alto/Tenor/Baritone), Triathlon, Golf.