

O. FELIX AMORUWA

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

ML systems engineer and NeurIPS-published researcher with 20+ years of hands-on experience building complex ML systems — from hand-coded BPTT in C++ (2004) to a 12-algorithm RL post-training workbench benchmarking GRPO/DPO across TRL, VeRL, OpenRLHF, and NeMo RL with GPU Docker passthrough today. Built production LLM evaluation infrastructure (aeval) with adversarial safety testing, statistical rigor, and automated safety gates. Motivated by making AI safe and beneficial; eager to contribute empirical research output at Anthropic.

AI/ML RESEARCH & PROJECTS

RL Workbench — Post-Training RL 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), 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 algorithm-specific metric profiles, cross-tab workflow lineage tracking, and standardized throughput/memory/convergence benchmarking across frameworks — directly targeting the training and fine-tuning analysis challenges central to LLM alignment research.

aeval — AI Model Evaluation Platform (2025–2026)

- Built local-first LLM 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.
- Engineered statistical rigor into the evaluation pipeline: bootstrap confidence intervals, Welch's t-test, Cohen's d effect size, and saturation detection; CI/CD integration with regression detection and automated safety gates. Stack: FastAPI orchestrator, TimescaleDB, Redis job queue, Next.js dashboard, Ollama.

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); spans 413 to 8B parameters (19M-fold scale increase from original system).
- Original 2004 system: hand-coded neural network in C++ with custom BPTT implementation. 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 pipeline for robot model training using multimodal AI (Claude/GPT-4V) to score model outputs (grasp poses, segmentation maps, bounding boxes) against natural-language rubrics — reducing evaluation cycles from 72 hours to ~4 minutes.
- Zero-integration architecture captures from any visualization tool (RViz, Matplotlib, custom dashboards) via screen capture; generates structured PASS/FAIL reports with confidence scores and edge-case flagging.

Deep Learning Education Platform (2025–2026)

- Developed 60 interactive demos across 20 chapters covering Goodfellow's Deep Learning textbook — transformers, GANs, diffusion models, autoencoders, optimization.

Lawrence Berkeley National Laboratory

- Computational analysis of RNA/DNA sequences under Dr. Steven Holbrook; assessed secondary structure accuracy using Monte Carlo and AI algorithms.

PROFESSIONAL EXPERIENCE

STREAMIO AI | Oakland, CA

Founder & CEO — ML Systems & Multi-Agent Infrastructure (macOS, Linux, iOS & Web) 09/24 – Present

- Built OpenClaw multi-agent orchestration framework with gateway protocol, subagent delegation, profile management, and session switching — enabling coordinated AI agent workflows; integrated Claude via MCP SDK for real-time multimodal AI analysis of screen captures.
- Engineered real-time HLS livestreaming pipeline: multi-stream canvas compositing (up to 9 concurrent sources at 30fps), FFmpeg WebM-to-MPEG-TS transcoding, WebSocket communication layer, and local HTTP playback server — demonstrating complex systems engineering under operational reliability constraints.
- Built production Electron + React + TypeScript desktop application with 100+ components, Redux Toolkit state management, and native macOS/Linux support; shipped cross-platform builds (macOS DMG with code signing/notarization, Linux .deb) with native ScreenCaptureKit integration via Swift.
- Built end-to-end auth and payments pipeline using Kinde OAuth 2.0, Stripe subscriptions, and Electron SafeStorage; integrated ElevenLabs TTS/STT with 6 auto-classified voice profiles and SHA-256 audio caching.

FINTELLECT AI | Oakland, CA

Founder & CEO — LLM Orchestration & AI Agent Platform (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, structured output validation, and token budget optimization — evaluating and fine-tuning model behavior across providers.
- Built domain-specific conversational AI agents scoped to distinct financial focal points; integrated LLM models for automated trade analysis at scale with AI-powered charting and macroeconomic analysis tools.

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; scaled throughput from 6K to 50K TPS via rSocket migration supporting ~1.5M concurrent connections with sub-25ms TP99 — large-scale distributed systems at production scale.
- Worked closely with telemetry and usage data (SQL, BigQuery) to prioritize developer pain points across ~20 mobile apps and 30+ product SKUs; delivered ICE Self-Service platform reducing developer onboarding from 2–3 weeks to minutes.

- Extended Java and Python SDK Starter Kits with scaffolding templates, build configurations (Gradle/Maven), testing frameworks, and CI/CD integration; led Mailchimp GCP-to-AWS migration for MSaaS.

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); led query performance optimization achieving up to 10x improvements in SCS search for enterprise customers.
- Delivered Scheduler Service end-to-end in ~4 months, enabling scheduled search capabilities for first-party applications; demoed at Splunk .conf19.

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); built Redis and XC10 caching layer across the enterprise for scalability, fault tolerance, and data redundancy.

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 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.

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