

# Roger Flores

Redwood City, CA • rflores3113@gmail.com • (650) 785-1299

[roger-flores.vercel.app](#) • [github.com/RogerFlores3113](#) • [linkedin.com/in/roger-flores-3113-nu](#)

---

*AI engineer who builds end-to-end (model layer to product layer) and ships to production.*

## TECHNICAL SKILLS

**Languages:** Python, TypeScript, JavaScript, SQL, Node.js

**AI / ML:** LiteLLM, OpenAI & Anthropic APIs, RAG, vector search, embeddings, PyTorch, scikit-learn, Pandas, NumPy

**Backend:** FastAPI, Postgres/pgvector, SQLAlchemy, Alembic, Redis/arq, Elasticsearch

**Frontend:** React, Next.js, MUI, front-end development

**Cloud / DevOps:** AWS (ECS Fargate, RDS, KMS, S3, ElastiCache, ALB), Terraform, Docker, GitHub Actions, Git, build systems

---

## EXPERIENCE

**C3.AI** (via Paradyne Management)

Redwood City, CA

*AI/ML Software Engineer*

Jan 2025 – Aug 2025

- On a 6-engineer team, maintained a production inventory-optimization system for a billion-dollar client; owned bug fixes, front-end pages, and job-management work across C3's React/Python platform (Cassandra + SQL).
- Surge-staffed a multi-county property-appraisal rollout on a 40%-compressed timeline; ramped on a new codebase in under two weeks and shipped all assigned features on the revised deadline.
- Developed a software version upgrade for C3's largest customer, coordinating across engineering and client stakeholders to deliver the migration on time and without service disruption.

**BigHat Biosciences**

San Mateo, CA

*Software Engineer Intern*

May 2023 – Aug 2023

- Shipped end-to-end features on an internal lab inventory system (React, Python, Postgres, AWS), from ticket to deploy, working directly with scientists and product at startup pace.

**Genentech**

South San Francisco, CA

*AI/ML Engineer Co-op*

Jul 2022 – Dec 2022

- Trained and evaluated ML models across scientific domains on an HPC cluster via SLURM; built researcher-facing dashboards that surfaced actionable insights from experimental data.

**Northeastern University**

Boston, MA

*Teaching Assistant - Machine Learning 2*

Jan 2024 – Apr 2024

- Produced 20+ exam prep videos for the full cohort; class average rose 15% vs. the prior term.

---

## PROJECTS

**Podium** | Personal AI Assistant Platform [podium-beta.vercel.app](#)

*Python, FastAPI, Postgres/pgvector, Next.js, AWS, Terraform*

- Built solo: a production-grade, cloud-hosted multi-tenant AI assistant that lets users chat with their documents, run code, and search the web from a single interface, with document RAG, agentic tool use, long-term memory, and model-agnostic routing via LiteLLM.
- Deployed to AWS via Terraform (ECS Fargate, RDS, ElastiCache, S3, ALB) with GitHub Actions CI/CD, Clerk auth, per-user KMS-encrypted API keys, and Redis-backed rate limiting.
- Implemented an agent loop with a pluggable tool registry (Tavily web search, E2B-sandboxed Python, document retrieval) and a memory system that extracts facts from conversations and surfaces them in future prompts.

---

## EDUCATION

**Northeastern University** | Khoury Honors Program

Boston, MA

*B.S. Computer Science, Minor in Mathematics*

Sep 2020 – Aug 2024