SHANAY GAITONDE

+1 (925) 487-1746 ◇ San Francisco, CA
shanaygaitonde@gmail.com ◇ linkedin.com/in/shanay-gaitonde ◇ github.com/shanayg15

EDUCATION

Bachelor of Science in Artificial Intelligence, University of California, San Diego — Expected June 2029

SKILLS

Languages	Python, Java, SQL, JavaScript, TypeScript, Prolog, Kotlin
AI/ML	LangChain, Multi-Agent Systems, RAG, NLP, Computer Vision, CNNs, LSTMs
Frameworks & Tools	TensorFlow, PyTorch, PostgreSQL, FastAPI, Node.js, Docker, Git, React, Next.js, GCP

EXPERIENCE

Software Engineer — May 2026 - Present
OpenSwarm — *San Francisco, CA*

- Building a full-stack marketing automation platform (React, Python/FastAPI, Docker) that orchestrates video publishing across fleets of Instagram and TikTok accounts via Android device automation, with job queueing, engagement tracking, and performance analytics; previously shipped a Next.js/TypeScript video learning platform with PostgreSQL-backed semantic search (vector embeddings) for user onboarding.
- Developed 10 TypeScript applications for OpenSwarm's in-app collection, spanning dynamic pricing, AI content generation, and compliance tooling, extending the core platform to niche use cases.

Founder — Apr 2026 - Present
Optivia — *San Francisco, CA*

- Architecting a multi-agent orchestration engine (Python, Claude Agent SDK) that decomposes software projects into dependency-aware task graphs, staffs each task with a specialized context-scoped agent, and executes with automated test verification, self-correction on failure, and cost-controlled routing across model tiers.
- Building a desktop client (Electron, React, TypeScript) featuring an interactive execution-graph interface that visualizes task dependencies, agent fleets, and build status in real time, backed by local SQLite persistence.

AI Research Intern — Oct 2024 - Jan 2026
AIEA Lab, UC Santa Cruz — *Santa Cruz, CA*

- Deployed the lab's ProSLM language model platform to Google Cloud Platform alongside undergraduate and graduate researchers, unifying multiple LLM research components into a single deployable system.
- Completed research onboarding in neuro-symbolic AI, building LLM-to-Prolog reasoning pipelines (Python, OpenAI API, SWI-Prolog), reimplementing the Logic-LM paper, and writing a backward-chaining inference engine from scratch.

PROJECTS

SixthSense (2nd Place, Qualcomm x Meta Hackathon). Built an on-device navigation copilot for blind and low-vision users: an Android (Kotlin) app runs Depth-Anything-V2 and YOLOv11 in parallel via ExecuTorch on the Qualcomm QNN backend, fuses depth and object detections into per-zone guidance, and steers the wearer through a BLE haptic belt (ESP32) and an offline on-device LLM voice agent.

HealthFlow (1st Place, Scalekit x Apify Hackathon). Built a 9-agent emergency-care pipeline (TypeScript, Next.js, LangChain) taking paramedic voice dictation to physician-approved EHR orders in under 60 seconds: agents structure transcripts into FHIR R4 records, run differential diagnosis, screen drug contraindications, and block unsafe orders behind physician review, all recorded on a SHA-256 immutable audit chain.

SignSpeak (2nd Place, Alameda County Science and Engineering Fair). Built a wearable real-time American Sign Language translator on a Raspberry Pi with camera and speaker modules, training a MobileNetV2 CNN (95% accuracy) and BiLSTM sequence model (90% accuracy) to convert signing into spoken audio.