

SAHIEL BOSE

✉ sahielbose@ucla.edu | github.com/sahielbose | [linkedin.com/in/sahiel-bose](https://www.linkedin.com/in/sahiel-bose) | sahielbose.com

EDUCATION

University of California, Los Angeles

Los Angeles, CA

B.S. in Mathematics of Computation

EXPERIENCE

Johns Hopkins University

Baltimore, MD

Research Software Engineer

August 2026 - Present

- Integrating nnInteractive, a 3D interactive segmentation framework, into BodyMaps, Prof. Zongwei Zhou's medical-AI program (Johns Hopkins CS and Medicine), for in-browser CT structure segmentation.
- Fine-tuning the segmentation model on Johns Hopkins Medicine CT data to improve per-structure accuracy for the 3D body atlas built for early cancer detection.

OpenSwarm

San Francisco, CA

Founding Engineer

May 2026 - Present

- Developing OpenSwarm, an Entrepreneur First-backed agentic browser with 5k+ active users: built a system that automates social media across 5 platforms, handling scheduled posting and analytics.
- Author end-to-end agentic workflows on OpenSwarm spanning creative generation, media tagging, CRM automation, architecture mapping, and scheduled overnight jobs and briefings.

ThinkNeuro LLC

Sacramento, CA

Research Engineer

November 2024 - Present

- Leading ML integration for a nonprovisional patent-pending wearable for real-time opioid overdose detection; built GRU, LSTM, and logistic-regression classifiers on a 721K-sample dataset with subject-level splits.
- Engineered the biphasic-overdose feature set and co-authored the research paper and literature review; built the React Native app that streams live vitals off a four-sensor Arduino Nano ESP32 array and alerts caregivers.

UC Santa Cruz, AIEA Lab

Santa Cruz, CA

Artificial Intelligence Researcher

October 2024 - April 2026

- Completed a 10-month onboarding in LLM logic and reasoning (symbolic inference, AI explainability) across Python and Prolog, then joined active research on ProSLM, the lab's logic-augmented LLM platform.
- Built natural-language-to-Prolog pipelines in SWI-Prolog and a from-scratch backward-chaining inference engine, and designed chain-of-thought scaffolds with self-consistency to refine LLM reasoning.

UC Riverside Research Lab

Riverside, CA

Machine Learning Researcher

May 2025 - January 2026

- Assisting a PhD researcher on GNN-based traffic-congestion-hotspot detection across large-scale, heterogeneous road networks (intersection and event nodes), applying spatial analytics to real-time geospatial data.
- Trained and benchmarked graph-sampling methods (GraphSAGE, GraphSAINT, GRAPES, HMSG) and heterogeneous-graph architectures to improve scalability and predictive accuracy for congestion prediction.

PROJECTS

Optivia  | *Agentic Workflows, LangGraph*

March 2026 - August 2026

- Developed Optivia, a macOS desktop app that decomposes a project into a dependency-aware task graph and staffs each node with a scoped, model-tiered agent, running them in build order with verified context passed downstream.
- Built it end-to-end (Electron, React, Python/LangGraph engine over typed IPC) with the fleet running on a live streamable desktop, and verification gates that block any node that fails before it advances.

Vigil - Anthropic x Abridge  | *YOLO11-Pose, Voice Agents*

July 2026

- Named a **Finalist, top 6 of 115 teams** for a continuous ER waiting-room re-triage agent: a YOLO11-pose model tracks skeletons only for privacy, with fall and seizure detectors I tuned on real clips to cut false alarms.
- An ESI algorithm scores each patient's acuity at intake, then dynamically re-scores it against their chart as events fire, clamped so urgency only rises, before triggering a nurse call and an auto-written SOAP note and FHIR bundle.

SixthSense - Qualcomm x Meta  | *ExecuTorch, Computer Vision*

June 2026

- Won **2nd Place** for a wearable that runs four models on a Galaxy S25 camera feed, on-device via ExecuTorch (YOLOv11n, Depth-Anything-V2, Qwen2.5-0.5B-Instruct, ML Kit OCR), to guide blind and low-vision users.
- Built the vision-to-haptics link: a BLE client and the mapper that converts per-zone depth into directional motor cues, driving an Arduino ESP32 and vibration motors on the harness.

SignSpeak - ACSEF  | *CNNs, LSTMs*

December 2024 - March 2025

- Earned **2nd Place** at the Alameda County Science & Engineering Fair (ACSEF).
- Co-led an AI-powered ASL recognition system using MobileNetV2 transfer learning and a bidirectional LSTM, achieving 95% validation accuracy and deployed on Raspberry Pi for real-time inference.

TECHNICAL SKILLS

Agentic AI & LLM Systems: Multi-Agent Orchestration, LangGraph, LangChain, RAG, Voice Agents

Machine Learning & Perception: Python, PyTorch, scikit-learn, YOLO, OpenCV, GNNs, LSTM & GRU