

Stephen Hung

stephenhung@berkeley.edu | (909) 414-4364 | linkedin.com/in/stephen-h-hung | stephenhung.me | github.com/stephenhungg

Education

University of California, Berkeley

Expected: May 2027

Bachelor of Science in Electrical Engineering and Computer Sciences

- **Coursework:** Data Structures, Algorithms, Computer Architecture, Discrete Math, Computer Security, Artificial Intelligence
- **Activities:** Cal Theta Tau (Co-ed Professional Engineering Fraternity), Cal Blueprint (Nonprofit Technology Solutions Club)

Skills

Programming Languages: Python, JavaScript (ES6+), TypeScript, Java, C++, C, Ruby, SQL, MATLAB, HTML5/CSS3, Rust

Libraries/Frameworks: React.js/Native, Node.js, Express.js, FastAPI, Next.js, Tailwind CSS, LangChain, Expo, Ruby on Rails

Developer Tools: Git, GitHub, Postman, Docker, VS Code, Figma, Jest, Vercel, Render, Railway, Websockets, Linear, Travis CI

Databases & Cloud: MongoDB, Firebase, Supabase, AWS (S3, EC2, Lambda), GCP (Cloud Storage, Vertex AI), Redis, SQLite

Relevant Certifications: Stanford University Machine Learning Specialization, Google IT Automation with Python Specialization

Experience

Software Developer, Blueprint, Replate Food Rescue Platform | *Berkeley, CA*

Sept 2025 – Present

- Developing a mobile app using **React Native** and **Expo** to digitize food donation tracking for **20+** partner nonprofit orgs.
- Architected full authentication system with signup, login, account management, and comprehensive error handling using **React Native** forms and **React Context API**, extending Replate's existing **Ruby on Rails** backend with secure API endpoints.

Software Engineering Consultant, ClearPath Medical | *Tustin, CA (Remote)*

Sept 2025 – Dec 2025

- Developed **AI powered PFMEA automation tool** using **FastAPI** and **LLMs** with **self validating agentic pipeline** that converts work instruction PDFs into risk assessments, reducing analysis time for FDA compliant medical device QC by **85%**.
- Engineered security first architecture with **WebSocket real time tracking**, multi stage validation system (ANALYZE → RATE → VALIDATE → CORRECT), and fully local processing with **SQLite** storage, ensuring HIPAA compliant handling.

Software Engineering Intern, OptiGenix | *Berkeley, CA*

May 2025 – Aug 2025

- Achieved **92.3% extraction accuracy** on labeled test data by training a **generative AI** model using Google Vertex AI, structuring detailed blood marker data from **60 unstructured** blood test PDFs stored and processed in **GCP cloud storage**.
- Enabled secure and scalable ingestion of clinical data by designing and deploying **HIPAA compliant data workflows** on **Google Cloud Platform (GCP)**, supporting 7,500+ monthly PDF uploads with **Cloud Storage**, **IAM role based access control**, **server side encryption**, and comprehensive **audit logging and monitoring** for regulatory compliance/tracing.

Projects

yolodex (Agentic YOLO Training) - OpenAI Codex Hackathon Winner | *YOLOv8, Codex*

[Live Demo](#) 

- Built Codex-native agentic pipeline that transforms raw unlabeled video into trained **YOLOv8** object detection models, automating frame extraction, labeling, augmentation, training, and evaluation with configurable **mAP@50** accuracy targets.
- Engineered **parallel subagent orchestration** spawning configurable Codex subagents in isolated git worktrees for distributed data labeling, with an autonomous iterative train-eval loop that re-labels and retrains until target accuracy threshold is met.
- Presented on stage at **OpenAI HQ** to judges including **Sam Altman**, **Greg Brockman**, **Sonya Huang** (Sequoia), and **Lenny Rachitsky**; featured on OpenAI's official socials, generating **500K+ impressions** and **20K+ reactions** in 24 hours.

opal (AI Gaming Companion) - NexHacks @ CMU | *LiveKit, YOLOv8, Claude CUA*

[Devpost](#)  | [Live Demo](#) 

- Developed autonomous **AI gaming agent** with game agnostic architecture supporting both FPS titles and turn based strategy games, joins voice channels and plays as a real teammate with **400ms latency**, growing a community of **1000+ members**
- Engineered custom **YOLOv8 and depth estimation** vision model achieving **60+ FPS** enemy detection for FPS games, **Claude** for strategic decision making, **LiveKit WebRTC** for real time voice, and cloud gaming to remain anti cheat compliant.

flow (Spatial Learning Platform) - SB Hacks XII Winner | *ElevenLabs, Marble, Gemini*

[Devpost](#)  | [Live Demo](#) 

- Won **President's Pick** and **MLH Best Use of ElevenLabs** for spatial learning platform where users speak concepts and explore photorealistic 3D Gaussian Splat environments with AI narrated overlays and contextual Q&A powered by **Gemini**.
- Engineered **6 API async orchestration** with **WebSocket progress tracking** through 5 minute pipeline (Deepgram → Gemini → Marble → SparkJS) with scene caching, rate limiting, and user world library deployed on **Vultr/MongoDB Atlas**.

darwin (Multi Agent Coding) - Cal Hacks 12.0 Winner | *React, Letta, LiveKit*

[Devpost](#)  | [Live Demo](#) 

- Built **6 agent orchestration system** where 4 specialized AI coding agents generate UIs from natural language while a commentator agent analyzes outputs via screen capture with **Letta cross agent memory** for context sharing and live narration.
- Engineered real time platform integrating **LiveKit STT** for voice commands, **ElevenLabs TTS** for agent commentary, synchronized live chat, audio reactive **Three.js/WebGL** visualizations synced to voice, and gasless **Sui blockchain** tipping.

clarifai (AI Research Agent) - NVIDIA AI Hackathon | *FastAPI, LangChain, Docker*

[Live Demo](#) 

- Built research paper analysis tool that generates 3Blue1Brown style animated video explanations, code implementations, and interactive Q&A from uploaded academic papers using **agentic pipeline** with concept extraction and multi modal generation.
- Engineered **self correcting LangChain pipeline** that autonomously writes, executes, debugs, and stitches Manim animation code with **parallel rendering** across multiple processes and real time **WebSocket** progress tracking for video production.