

Saicharan Ramineni

904-615-0799 | csramineni@gmail.com | linkedin.com/in/saicharanramineni | github.com/Godlydonuts

EDUCATION

University of Central Florida

Bachelor of Science in Computer Science; GPA: 4.0

Orlando, FL

Expected May 2027

- Honors: Burnett Honors College, Dean's List.
- Involvement: Knight Hacks, UCF Competitive Programming Team, Society of Asian Scientists and Engineers (SASE), Indian Student Association (ISA).

PROJECTS

Project Sentinel | *React, Cerebras, Gemini, Deepgram, ElevenLabs, WebSockets, WorkOS* Nov. 2025

- **Awarded Grand Prize (3rd Place)** at the Liquid Metals AI Championship, an international hackathon, for engineering an enterprise-grade AI guardian against voice phishing.
- Architected a high-speed AI pipeline using Deepgram Nova-2 for STT and Cerebras (Llama 3.3-70b) for **sub-second semantic threat analysis** of live conversation patterns.
- Engineered an automated response system using ElevenLabs Turbo v2 to generate ultra-realistic voice warnings with **sub-300ms latency** to alert users of social engineering threats.
- Built a high-performance React frontend with 3D visualizations using Three.js and native WebSockets for **full-duplex, low-latency** audio streaming.

DevDraft AI | *Python, FastAPI, Apache Kafka, Confluent Cloud, Google Cloud Provider, Vertex AI* Jan. 2026

- Engineered a multi-modal agentic "AI Senior PM" that converts live meeting audio into **production-ready React + Vite codebases** in real-time.
- Architected an **event-driven AI pipeline** using Apache Kafka and Confluent Cloud to orchestrate audio streams and transcript processing across decoupled microservices.
- Developed "Concept Graph" logic to handle **real-time user pivots**, allowing the system to instantly update project PRDs and strike through outdated requirements mid-conversation.
- Integrated an automated coding engine using Gemini 3.0 Pro/Flash to **parallelize file generation** and export optimized prompts for agentic tools like Replit Agent and Bolt.new.

Vectorized-2048 | *JAX, Flax, Python, Reinforcement Learning, Deep Q-Networks* Feb. 2026

- Engineered a massively parallel Deep Reinforcement Learning environment for 2048 in JAX, achieving a **4,000x speedup** (~2,000,000 steps/sec) over standard sequential Python implementations.
- Architected the game engine as **stateless, branchless** matrix transformations, eliminating CPU bottlenecks by running the entire simulation and training loop end-to-end on the GPU/TPU.
- Implemented a **Dueling Double DQN (DQN)** with learned tile embeddings and custom convolutional kernels to capture complex row/column sliding patterns and improve sample efficiency.
- Optimized training stability through logarithmic reward scaling and gradient clipping, enabling the agent to **simulate 4,096 games simultaneously** as a single batched operation.

CityMate | *React, TypeScript, Python, Flask, Scikit-learn, Supabase, Selenium* Mar. 2025

- Developed an intelligent real estate platform using a Scikit-learn ML model to analyze interaction patterns and deliver **personalized recommendations from 10,000+ listings**.
- Automated a robust data ingestion pipeline using Selenium and BeautifulSoup to aggregate property data and enrich listings with **Google Places API** commute/amenity metrics.
- Implemented **real-time database capabilities** via Supabase for instantaneous listing updates and integrated agent-user chat functionality.

TECHNICAL SKILLS

Languages: Python, TypeScript, JavaScript, Java, C, SQL, Dart, LabVIEW

AI & Machine Learning: JAX, Flax, Reinforcement Learning (DQN), Vertex AI, Gemini, Cerebras, Scikit-learn, Deepgram, ElevenLabs

Frameworks & Libraries: React 19, FastAPI, Flask, Flutter, Tailwind CSS, Framer Motion, Three.js, Node.js, Pandas, NumPy, Matplotlib

Tools & Infrastructure: Apache Kafka, Confluent Cloud, AWS (S3, Lambda, EC2), Google Cloud Platform, Supabase, Docker, Git, Kubernetes, Selenium, WorkOS, Stripe