

Luan Luong

El Monte, CA | [LinkedIn](#) | 626-297-3685 | luongryanll@gmail.com

SUMMARY

CS junior at Cal State Fullerton building full-stack AI applications in production. Passionate about developer tooling, automation, and using modern software to solve real engineering problems at scale

EDUCATION

California State University Fullerton
BS in Computer Science

Fullerton, CA
Graduation Date: June 2027

SKILLS & TECHNOLOGIES

Languages: C++, Python, JavaScript, TypeScript

Web & Frameworks: HTML, CSS, React, Node.js, Next.js, FastAPI, REST APIs

Databases: SQL Server, PostgreSQL

Cloud & Tools: AWS (Textract, S3), Microsoft Azure, Git/GitHub, VS Code

PROJECT EXPERIENCE

Pitchside — Fantasy Football Agent | Google Cloud Hackathon | [Source](#) / [Demo](#)

- Built an agentic AI system using Google ADK and Gemini 2.5 Flash on Vertex AI with 15 custom tools — lineup optimization, match prediction via a custom probability model, and live standings from an external API — backed by MongoDB Atlas seeded with 3,230 match statistics from the 2022 World Cup
- Engineered a full-stack fantasy scoring engine in Python/FastAPI with a Next.js frontend featuring a real-time interactive pitch view that updates live when the agent recommends a lineup, deployed on Railway and Vercel

MedVoice — Powered Healthcare Bill Navigator | [Source](#) / [Demo](#)

- Built a full-stack bill navigator using FastAPI, Next.js, and PostgreSQL, integrating AWS Textract for OCR and Claude API for multilingual document analysis across 8 languages with CI/CD via GitHub Actions
- Implemented Clerk authentication with Google OAuth, a personal analytics dashboard using Recharts displaying spending over time, most common charges, and bill history with Motion animations and full mobile responsiveness
- Built an interactive provider search page using the Google Maps and Places API with real-time ratings, photos, and a drag-to-resize split panel layout, and added AI-powered appeal letter generation and a contextual bill chat feature

PodMind — English Learning Platform via Real Podcasts | [Source](#) / [Demo](#)

- Built a full-stack podcast learning platform using Node.js, Express, Next.js, and PostgreSQL, implementing an end-to-end audio pipeline: RSS ingestion → ffmpeg compression → OpenAI Whisper transcription → pgvector embeddings for semantic search across episode transcripts
- Engineered a synchronized audio-transcript player with word-level timestamps, click-to-translate via DeepL + Claude API, and a spaced repetition vocabulary system (SM-2 algorithm) — enabling immersive language learning directly from native English content

Mini-SQL (C++)

- Built a C++ SQL interpreter with a full lexer → parser → AST → executor pipeline supporting SELECT, INSERT, UPDATE, DELETE with WHERE filtering and column projection on in-memory tables