

AYAAN KHAN

(773)-937-2448 . ayaanahmedkhan12@gmail.com . in/ayaanahmedkhan/ . github.com/ayaan47-1

EDUCATION

ILLINOIS INSTITUTE OF TECHNOLOGY | Chicago, Illinois
Bachelor of Science in Artificial Intelligence; **Minor** in Architecture

Graduating May 2026

Relevant Coursework: DSA, AI, ML, NLP, DBMS, Assembly, Data Mining, Discrete Math, Linear Algebra, Probability, Statistics, OOP

SUMMARY

AI engineer with two years at a construction tech firm, shipping production RAG systems and agent workflows across real projects. Built and deployed retrieval pipelines, async backends, and agentic BIM tools used by actual project teams. Comfortable owning backend-to-deployment: FastAPI, PostgreSQL/pgvector, Redis, Docker, LiteLLM, Inngest.

WORK EXPERIENCE

D.A. SYNTEC LTD. | Chicago, Illinois

- AI & Digital Developer (Returning Intern)

January 2026 - Present

- Built and deployed a production RAG chatbot over firm documents -- chunking, embedding, indexing with ChromaDB -- now used by internal teams to answer project queries without digging through file shares.
- Built SyntecAgent: natural language CRUD over a BIM classification database, with atomic SQLite to ChromaDB sync and confirmation gates before any destructive writes. Used by project teams in production.

- AI & Digital Developer (Intern)

May 2025 - December 2025

- Curated and published weekly posts using Canva across LinkedIn and Mailchimp; drove over 60% growth in impressions and a notable increase in audience engagement following a content strategy restructure.
- Built a custom RAG pipeline using OpenAI and ChromaDB -- chunking, embedding, and indexing firm documents to enable natural language retrieval -- serving as the technical foundation for later production chatbot deployment.

PROJECTS

SYNTEC AGENT | Professional Project

February 2026 - Present

- Shipped a semantic search and retrieval augmented system using GPT 4o and ChromaDB embeddings to deliver question answering across PDF documents, CSVs, blog posts, and website content with automated source citations.
- Built an agentic BIM classification system using OpenAI function calling that enables natural language CRUD on coding, classification, and naming entries, with atomic SQLite to ChromaDB sync and confirmation gates on destructive actions.
- Containerized the full stack with Docker, deployed the frontend to Vercel with HTTPS and server side API rewrites, and added Redis caching with a DeepSeek fallback LLM for cost optimization.

VERITASLAYER | Founder and Solo Developer

March 2026 - Present

- Solo built and deployed a 13 stage AI pipeline (FastAPI, Inngest, PostgreSQL, LiteLLM) that extracts obligations, risks, and deadlines from construction contracts and CRE leases -- every claim backed by a verbatim quote, page number, and confidence score.
- Built a production Next.js frontend with Clerk auth, human review workflows, real-time status polling, and an eval harness (Jaccard quote matching for precision/recall) used to diagnose and fix chunk selection accuracy.

CLUTCH | Co-Founder and Developer

January 2026 - Present

- Built Clutch: a course generation SaaS with a 5 stage agent pipeline, persisted job state, retry policies, and SSE streamed progress. Cut inference cost with LiteLLM routing and shared source deduplication.
- Implemented an asynchronous FastAPI backend with PostgreSQL plus pgvector and Redis caching, and streamed real time job progress to clients using SSE backed by Redis to keep generation workflows responsive.
- Cut inference cost and improved output quality using LiteLLM routing and shared source deduplication; enforced structured outputs with PydanticAI validation and added observability via Sentry and PostHog.

SPORTS AND METRICS TRACKER | Personal Project

January 2026

- Built an end to end soccer video analysis pipeline using YOLOv8 detection and ByteTrack tracking, adding temporal smoothing for stable motion analytics and CPU only local processing with robust cross-platform video input and output.
- Developed export and visualization tooling, including CSV and JSON metrics with track filtering (minimum 15 frames), heatmap rendering with percentile transforms, and tiered error handling with debug mode progress reporting.

SENTIMENT ANALYZER | Academic Project

September 2025 - December 2025

- Engineered a machine learning classifier using Logistic Regression with TF-IDF features (1-2 grams) to analyze over 50,000 Amazon fine food reviews, achieving high accuracy through L2 regularization and balanced class weighting on preprocessed text data.
- Built an interactive Streamlit web application with real time sentiment prediction and confidence scoring, implementing complete ML pipeline from data preprocessing (noise removal, text cleaning) to model deployment.

SKILLS

Languages: Python, SQL, Java, JavaScript, TypeScript, HTML/CSS

Frameworks & Tools: FastAPI, Flask, React, Svelte, Vector Databases, Redis, LiteLLM, PydanticAI, Inngest

ML and LLM: RAG, embeddings, APIs, LangChain, LangGraph, scikit-learn, PyTorch, Hugging Face

Infra: Docker, Git, CI/CD, Sentry, PostHog, Vercel