

- LLM Engineering
- Agentic AI
- RAG Systems
- Solution Architecture

Engineer with 6 years shipping products, the last 5 at IBM embedded with Fortune 500 clients designing and delivering custom AI systems end to end. Patent-pending work on LLM reliability and a track record of custom RAG and agentic systems shipped into production.

## Work Experience

### AI Customer Success Engineer

Jan 2025 — Present

IBM New Orleans, LA

- Embedded with account teams as the technical SME on IBM's enterprise LLM and agentic platforms (Watsonx.ai, Watsonx Orchestrate) and coding assistant (IBM Bob), scoping client AI problems and designing the solution in client conversations.
- Architect solutions and ship custom prototypes in Python, LangChain, and FastAPI that move client conversations past feasibility and into deployment.
- Carry technical ownership past deployment to evolve the system and unlock new use cases that expand the product footprint within the client.
- Lead a small team running multi-client in-person workshops built around the technical AI use cases attendees bring in.

### AI Engineer, Client Engineering

July 2021 — Dec 2024

IBM Boston, MA

- Led 15+ AI engagements end to end across Fortune 500 enterprise clients, embedded with client teams to design and ship custom AI systems in a fast-moving startup-like environment within IBM (Python, LangChain, FastAPI, Pinecone, ChromaDB, Watsonx).
- Invented a patent-pending approach for LLM-powered diagram generation on Watsonx (Llama, Mistral), improving reliability from 10% to 99% by separating creative LLM reasoning from deterministic code generation. Turned a failed project into a reusable organizational asset.
- Built a custom agentic RAG pipeline for a major fintech provider's API documentation with custom ingestion, chunking, and hybrid search with reranking, reducing support agent answer time from hours to minutes. Selected to demo at the client's annual conference in Las Vegas.
- Shipped a range of AI systems across generative workflows, document understanding pipelines, and classification systems, weighing architectural decisions on each to land on the right pattern, model, and data pipelines.
- Drove discovery and whiteboarding sessions with each client to pick the approach with the strongest ROI before any code was written.

## Projects

### AI Video Classifier and Recipe System

Nov 2024

yt-search-ui.vercel.app

- Built an end-to-end system to classify and organize 2,500+ YouTube videos using a fine-tuned OpenAI model trained on 250 hand-labeled examples, running daily on self-hosted infrastructure.
- Engineered hybrid search (keyword + semantic) over video embeddings in Pinecone, improving retrieval accuracy by 4x through a custom preprocessing pipeline that generates enhanced descriptions from transcripts before embedding.

### editsbycb.com — Personal Stylist Booking Platform

Dec 2024 — Present

editsbycb.com

- Built a full-stack commercial platform on Next.js, PostgreSQL, and Drizzle ORM with Stripe subscriptions, deployed on Vercel across dev, staging, and production environments.
- Owned the entire product end to end including pricing, UX, copy, and booking workflows, acquiring 15 paying customers on recurring revenue.

## Education

### Bachelor of Science, Computer Engineering — Minor in Computer Science

Graduated May 2021

Louisiana State University

Baton Rouge, LA

## Most Relevant Skills

LLM Application Development • Agentic AI & Multi-Agent Systems • RAG Systems • Python & FastAPI • Vector Databases (Pinecone, ChromaDB) • Prompt Engineering • Solution Architecture