

PARI GOYAL

Agentic AI Engineer | LLM Pipelines and Automation in Production

+91-6284013296 | goyalpari70@gmail.com | India | Open to relocation and remote

Portfolio | LinkedIn | GitHub

SUMMARY

I am the **only engineer** at a capital advisory firm, so the LLM pipeline running their deal flow is mine end to end. It has read and qualified over 200 startup leads so far. Nothing it writes reaches a client without someone approving it first. Most of my time goes on the part nobody puts in a demo, which is what the system does when the model is wrong and whether the next person can run it without me.

SKILLS

AI and Agentic Systems: LLM application development, Claude API, OpenAI API, prompt engineering, tool calling and function calling, agent workflow design, multi-step orchestration, structured output generation, prompt evaluation and test harnesses, human-in-the-loop approval patterns, AI guardrails, pipeline logging and observability, n8n, Claude Code, retrieval-augmented generation (RAG), semantic chunking, embeddings and vector retrieval, LangChain, LangGraph, FastAPI

Backend and Data: Python, TypeScript, Node.js, REST APIs, webhooks, PostgreSQL (Supabase), pgvector and vector search, MongoDB, SQL, row-level security, Stripe

Frontend: Next.js, React, TypeScript, Tailwind CSS, Vite

Infrastructure and Tooling: Vercel, Docker, Git and GitHub, JSON-LD structured data, Hunter.io

CS Fundamentals: Data Structures and Algorithms, OOP, DBMS, Operating Systems, Computer Networks

EXPERIENCE

Technology and Operations Associate (AI and Software Engineering), Akro Ventures | Apr 2026 - Present

- Built the deal-management platform from **nothing to production in three months**. It runs the firm's whole deal flow now, from the public intake form through AI analysis to the outreach that actually goes out the door.
- Built the LLM layer on the Claude API, using structured tool calls so raw intake data becomes deal summaries and hot / warm / cold qualification. **200+ startup leads** processed to date, and it is the system the team ran through an **INR 20 Cr** client fundraising mandate.
- Deals were approving themselves. Corporate email scanners were pre-fetching the approval links to check them for phishing, which counted as a click. Moved approval off the GET request onto an authenticated POST.
- The model drafts, it never sends. Approvals run on single-use tokens, and the factual claims about the firm live in application code where the model cannot reach them.
- Built an n8n pipeline over daily news sources that turned **roughly 10 hours a week** of content production into review and post.

Agentic AI Intern, Oyelabs | Feb - Apr 2026

- Built production AI agents for client projects on n8n, Supabase, and the Claude and OpenAI APIs, forming pipelines that classify, draft, and route work to a human before anything reaches a customer.
- Wrote an evaluation harness for the client-facing agents. Every workflow had to pass SOP cross-checks before it was allowed to run, which caught bad output before a customer saw it and gave the manager back **a full day a week** they had been spending on manual checks.

PROJECTS

RoleFit, Retrieval Over the Job Descriptions I Applied To | Live at rolefit-wine.vercel.app | FastAPI, LangGraph, pgvector

- I kept getting skill-mismatch rejections without being told which skill, so I put the job descriptions into Postgres and asked them. Retrieval is hybrid: pgvector embeddings for meaning, Postgres full-text for rare exact tokens like a framework name, fused by reciprocal rank fusion so two incomparable scores never need a weight between them.
- A LangGraph state machine grades its own retrieval, rewrites the query when the chunks come back thin, and says the corpus does not cover it rather than inventing an answer. Deployed on Vercel behind FastAPI, with an eval suite gating on deterministic citation validity, because an LLM judge scored my correct answers 0 out of 4 until I found the harness bug feeding it the wrong input.

AI Job Search Outreach Platform | Next.js, Claude API, Supabase

- I was losing **4 to 6 hours a day** to manual cold outreach, so I built something to do it. Company research and drafting now take **under 10 minutes** a day.
- Facts about my background are locked in code, so the model writes each pitch but cannot invent a credential. Addresses are verified through Hunter.io and an MX check before anything sends, follow-ups fire on day five on their own, and a dashboard breaks the results down by industry, city and channel.

AI Lead Qualification and Quotation Pipeline | n8n, Supabase, Stripe, LLM API

- A form submission kicks off email verification, a Supabase check for duplicates, then an AI pass that scores the lead hot, warm or cold. Every step logs so somebody else can pick it up.
- Quotations get drafted by the model and held at a manager approval gate before Stripe bills anyone. A human confirms before money moves.

EDUCATION

B.E. Computer Science and Engineering, Thapar Institute of Engineering and Technology (TIET), Patiala | 2022 - 2026