

ATHARV PATOLE

Jersey City, NJ | 201-800-1948 | atharvpatole6@gmail.com | [LinkedIn](#) | [GitHub](#)

SUMMARY

I build LLM systems that go all the way out into production, handling ingestion, retrieval, agents, and the frontend that someone actually opens. Over the last four years, I have specialized in turning messy public and geospatial data into actionable intelligence. I work directly with the people experiencing the problem, figuring out what they actually need versus what they ask for, and I put strict evaluations in place before shipping rather than after something breaks.

WORK EXPERIENCE

Fulton Ring | Data Scientist, AI Systems | New York, NY

Aug 2025 - Jul 2026

- Analysts at a regional housing nonprofit queued behind engineers for every data pull. Shipped a production LLM search layer over housing, demographic, and environmental datasets so they could ask in plain language, cutting manual review time roughly 40%.
- Retrieval kept surfacing plausible but wrong parcels because embeddings ignored the identifiers analysts reason with. Rebuilt the RAG layer over Pinecone around structured metadata for property IDs, zoning codes, and spatial indicators, lifting retrieval consistency close to 30%.
- Groundedness checks before each release took hours of reading outputs by hand. Instrumented MCP tool-calling with LangSmith tracing and custom LLM-as-judge guardrails, then gated releases on regression evals, compressing that review to minutes.
- Open-data records arrived inconsistent enough that the program team re-cleaned them every cycle. Automated cross-source validation and anomaly detection across millions of records via Python and SQL pipelines on AWS, with match rules agreed alongside the data owners.
- Architects, engineers, and investors had no route into the property data without borrowing a data team. Launched standalone Next.js and MapLibre sites reading spatial files above 10GB, trimming query latency roughly 60%.

Scon Infra | Analytics Engineer | Mumbai, India

Jul 2020 - Jun 2023

- Finance and engineering leads reconciled infrastructure planning figures by hand each month, stalling sign-off. Scripted the reporting logic across more than \$12M in planning workflows and reduced recurring manual effort about 35%.
- Downstream models kept breaking on inconsistent source data. Wrote Python and SQL ingestion pipelines for real estate and construction datasets, settling validation rules with the engineering team, which retired the manual prep step entirely.
- Executives waited on hand-assembled reports before committing. Defined governance rules and relational models spanning engineering and finance, shortening decision turnaround nearly 45%.

PROJECTS

SpiderNet | Geospatial AI Search Platform | LLM Agents, RAG, Pinecone, Python, Next.js, AWS | [App Link](#)

- Solo built and deployed a platform indexing hundreds of thousands of municipal property records in Pinecone, resolving natural-language land-use questions through LLM agents. Roughly 20 people use it.
- An initial Kriging and XGBoost model failed on sensor-siting bias, so moved to a physics-based HAND-stage inundation method, validated across 241 sensors against NYC DEP's Stormwater Map and 311 flood complaints, then integrated the resulting risk layer over roughly 1 million NYC properties.

CareFlow Copilot | LLM Healthcare Analytics Assistant | Python, Next.js, Cloud Deployment | [App Link](#)

- Deployed an assistant answering natural-language questions over hospital encounter, utilization, and time-series data, with tool-calling workflows that turn raw records into structured metrics and plain-English summaries for non-technical staff.

TECHNICAL SKILLS

Agentic AI and LLM Systems: LLM agents, multi-step workflows, tool and function calling, LangGraph, OpenAI Agents SDK, Model Context Protocol (MCP), prompt and context engineering, human-in-the-loop workflows

RAG and Retrieval: Graph RAG, LlamaIndex, hybrid retrieval (BM25 and dense), semantic chunking, semantic caching, knowledge graphs, Pinecone, Chroma, Vertex AI Vector Search

Evaluation and Observability: LangSmith, Langfuse, OpenTelemetry, LLM-as-judge, regression evals, cost-per-query tracking, prompt-injection defense, PII filtering

Languages and Backend: Python, SQL, TypeScript, JavaScript, Flask, REST APIs, Next.js, Vercel, MapLibre

Data and Cloud: AWS, GCP, ETL and ELT pipelines, PySpark, Docker, Kubernetes, PostgreSQL (Supabase), SQL Server, MongoDB, DuckDB

Machine Learning: Classification, clustering, time series forecasting, anomaly detection, feature engineering, model evaluation

EDUCATION

Pace University | M.S. Computer Science | New York, NY

May 2025

Sanjay Ghodawat University | B.Tech Civil Engineering India

May 2021