

ADITHYA SINGH CHAUHAN

DataOps · AI Systems & Backend Platforms · Data Engineering

Bengaluru, India | 6309464585 | adithyachauhan4@gmail.com

linkedin.com/in/AdithyaChauhan | [GitHub](https://github.com) | [Portfolio](https://portfolio)

SUMMARY

DataOps Intern with hands-on experience across cloud platforms, DevOps tooling, and AI application development. Completed a structured 6-month technical program spanning Azure, AWS, Docker, Kubernetes, Terraform, Databricks, and GenAI tooling.

Designed and shipped a production-grade multi-agent LLM application as a 24-day capstone project at Sigmoid Analytics, using FastAPI, LangGraph, and PostgreSQL with end-to-end observability, automated evaluation, and CI/CD. Strengths in system design, data modeling, state management, and building observable, production-ready services.

EXPERIENCE

DataOps Intern — Sigmoid Analytics

November 2025 – June 2026

Python, FastAPI, LangGraph, GPT-4o-mini, PostgreSQL, SQLAlchemy, Docker, Kubernetes, Prometheus, Grafana, LangSmith

- Completed a 6-month cross-domain technical training program covering cloud platforms (Azure, AWS), DevOps tooling (Docker, Kubernetes, Helm, Terraform), data platforms (Databricks, Hadoop), and GenAI development (LangChain, FastAPI), with assessments.
- Designed and built an AI-powered customer support platform as a 24-day capstone project, spanning product discovery, order management, and support workflows through a multi-agent architecture — an LLM router classifying intent across 3 specialist agents (7 compiled state graphs) sharing one typed state object.
- Architected deterministic, security-critical business logic for search, ownership verification, and severity assessment, confining the LLM to language tasks to keep correctness and access control (IDOR-safe) outside model behavior.
- Built session-aware state management with query-relaxation strategies and an 8-signal weighted ranking model across a 1,300+ product catalog backed by an 8-table PostgreSQL schema.
- Established automated quality gates with unit testing, DeepEval-based LLM evaluation, and an 80% coverage threshold integrated into a 6-stage CI/CD pipeline (lint, test, Docker build, Helm deploy, evaluation).
- Instrumented full observability — Prometheus metrics, LangSmith tracing with per-turn feedback, and correlation-ID logging to Loki/Grafana — across containerized services deployed via Docker and Helm.

PROJECTS

Azure Databricks Real Estate ETL Pipeline

Self-Directed · Applied DataOps Coursework

Azure Databricks, PySpark, Delta Lake, ADLS Gen2, Databricks Workflows, SQL

- Developed an end-to-end Medallion architecture (Bronze → Silver → Gold) pipeline processing real-estate transactions, units, and projects into 3 Gold KPI tables covering sales trends, area performance, and supply distribution.
- Implemented PySpark Silver-layer transformations (type casting, null handling, regex cleaning, multi-source joins) persisted as Delta Lake tables.
- Orchestrated the full pipeline with Databricks Workflows and surfaced KPIs via a SQL dashboard with a Power BI bridge.

Azure Data Factory Incremental Ingestion Pipeline

Self-Directed · Applied DataOps Coursework

Azure Data Factory, Azure SQL Database, ADLS Gen2, Parquet

- Designed a metadata-driven, watermark-based incremental ingestion pipeline (Lookup → If-Condition → Copy) loading only changed rows from Azure SQL into an ADLS Gen2 Bronze layer.
- Parameterized source queries by schema, table, and CDC column, reducing per-table pipeline code to a single reusable template.
- Wrapped the parameterized pipeline in a ForEach loop to ingest multiple tables into partitioned Parquet storage, with GitHub-based version control.

Full-Stack E-Commerce Platform (MERN)

Personal Project

MongoDB, Express.js, React, Node.js, Stripe, JWT

- Shipped a full-stack e-commerce app with separate customer and admin frontends, supporting 50+ products and secure checkout.
- Implemented role-based access control with JWT and bcrypt, and integrated Stripe payments for end-to-end order processing.
- Built category-based filtering that cut product search time by ~30% and optimized MongoDB queries for ~20% faster product loads.

TECHNICAL SKILLS

Core: Python, SQL, JavaScript, TypeScript, FastAPI, PostgreSQL, SQLAlchemy, REST APIs, System Design, Data Modeling

AI Systems: LangGraph, LangChain, Agentic AI, LLM Applications, State Management, Evaluation (DeepEval), Production Systems

Data & Cloud: Azure, AWS, Azure Data Factory, Databricks, Delta Lake, ADLS Gen2, PySpark, ETL, Medallion Architecture

DevOps: Docker, Kubernetes, Helm, Terraform, Azure DevOps, GitHub Actions, CI/CD

Observability: Prometheus, Grafana, Loki, LangSmith

Frontend: React, Next.js, Tailwind CSS

EDUCATION

B.Tech, Electronics and Communication Engineering

2024

National Institute of Technology, Srinagar