

SHREYA BHAT

GenAI Engineer | LLMs | RAG | AI Applications

Bangalore, India | shreyabhat545@gmail.com | +91 9380040180 | [LinkedIn](#) | [GitHub](#) | [Portfolio](#)

EXPERIENCE

Data Science & Analytics Trainee | | [CIPLA Ltd](#)

July 2024 - June 2025

- Responsible for building RAG-based document-processing system using LangChain + pgvector to index 1,000+ equipment logs & SOPs, which resulted in a reduction of monthly trouble-shooting job cards from 120+ to approximately 40 (about 67%).
- Developed and applied prompt engineering to optimize the retrieval and response quality of each of the multiple LLM-backed workflows.
- Assessed the RAG pipeline output with a 100-observation golden dataset via RAGAS metrics (i.e., faithfulness, context recall and answer relevancy).
- Designed and implemented an anomaly detection pipeline across 1,400+ sensors (15 SCADA and 47 BMS systems) and linked this anomaly detection pipeline with a knowledge retrieval system; thereby decreasing the time required for root-cause analysis by 55-60%.

KEY PROJECTS

AskQL – Natural Language to SQL | [Link](#)

- Problem: Non-technical users need a way to query databases without writing SQL.
- Solution: Built a natural-language-to-SQL pipeline using intent/entity extraction, rule-based parsing and LLM-based query generation, with FastAPI backend and React UI.
- Depth: Designed the pipeline to separate query understanding, SQL generation and execution, making it easier to extend and evaluate.

AI Hallucination Detection System | [Link](#)

- Problem: LLM responses can contain claims that sound plausible but lack supporting evidence.
- Solution: Built a multi-stage pipeline for claim extraction, web evidence retrieval, semantic similarity, using Gemini, FAISS and sentence embeddings.
- Depth: Added evaluation/trust scoring and achieved ~87% detection accuracy, exposing the system through FastAPI + React Dashboard

TruthLens - Fake News detection: MLOps | [Link](#)

- Problem: Fake-news classification needs to be more than an offline ML model; it needs a reproducible, testable and deployable pipeline.
- Solution: Built an end-to-end ML pipeline covering data validation, preprocessing, model training, evaluation and inference, with production-oriented project structure.
- Depth: Added automated testing, artifact/model management and deployment-ready APIs to make the model reproducible and maintainable.

SKILLS

Languages: Python, SQL

AI/ML: Machine Learning, NLP, Scikit-learn, XGBoost, PyTorch

Generative AI: LLMs, RAG, Prompt Engineering, LangChain, LangGraph, Hugging Face

AI Engineering: FastAPI, REST APIs, Vector Databases, FAISS, LLM Evaluation

Cloud & Tools: AWS, Docker, PostgreSQL, Git/GitHub

EDUCATION

PGP in Data Science (Specialization: Gen AI)

July 2025 - May 2026

Great Learning

Bachelors of Engineering

Nov 2020 - May 2024

Bangalore Institute of Technology, Bengaluru

ACHIEVEMENTS

- "Performer of the Month" – awarded 3 times at CIPLA Ltd for consistent, high-impact project delivery and stakeholder satisfaction.
- Developed end-to-end data and AI systems independently, from problem scoping to deployment and stakeholder presentation.