

Lokesh Kumar

+91-9302335988 | hellokumarlokesh@gmail.com | GitHub | LinkedIn | Portfolio

Raipur, Chhattisgarh

TECHNICAL SKILLS

AI/ML & LLM: LLM Application Development, RAG Systems, Prompt Engineering & Optimization, AI Agents, Embeddings, Vector Search, Generative AI, NLP, Model Deployment

Frameworks & Libraries: LangChain, HuggingFace, LlamaIndex, OpenAI SDK, FastAPI, Streamlit, TensorFlow, Scikit-learn, Pandas, NumPy

Vector Databases & Search: FAISS, BM25 (hybrid retrieval); exposure to Pinecone/Qdrant architectures

Programming & APIs: Python, JavaScript (Node.js), SQL, REST API Design, JSON, HTML/CSS

Cloud & DevOps: AWS (Cloud Practitioner + Solutions Architect Certified), Docker, Git, GitHub Actions, CI/CD

Data & Evaluation: ETL Pipelines, Data Augmentation, LLM Evaluation (LangSmith, Ragas, DeepEval), Structured & Unstructured Datasets

EXPERIENCE

•AI Engineer Internship

June – August 2025

Apana Time

Online

- Built and deployed an AI-powered Data Analyst Bot using FastAPI, Streamlit, LangChain, and Groq-hosted Llama LLMs – prototyped RAG pipelines, integrated FAISS vector search, and optimized prompts for quality, latency, and cost.
- Developed LLM-based AI agents with tool-use capabilities: SQL Query Generation, contextual RAG, automated data augmentation, session-based memory, and multi-format file support (CSV, PDF, images).
- Implemented LLM evaluation pipelines using LangSmith and Ragas for automated quality testing; containerized full deployment with Docker and validated REST APIs via Postman.
- Applied AI-first development practices including prompt engineering, LLM-as-copilot workflows, and structured SDLC – achieving measurable efficiency gains in specification writing and analytical automation.

PROJECTS

•DataAnalystBot – LLM-Powered Data Analysis Agent

AI agent for end-to-end data analysis, SQL generation, and RAG-based Q&A over user-uploaded datasets.

GitHub

- Built multi-tab Streamlit app backed by FastAPI; integrated Groq-hosted Llama LLMs with LangChain for SQL generation, contextual Q&A, and automated data augmentation (KNN imputation, IQR winsorization, synthetic row generation).
- Designed FAISS + HuggingFace embedding pipeline for semantic document retrieval; optimized prompts at low temperature for structured output accuracy; resolved Dockerization dependency conflicts for reliable containerized deployment.
- Tech: Python, FastAPI, LangChain, FAISS, HuggingFace, Groq (Llama), Streamlit, Docker.

•Enterprise RAG System

Scalable Retrieval-Augmented Generation system for secure multi-tenant document intelligence.

GitHub

- Engineered hybrid retrieval pipeline combining FAISS dense vector search with BM25 ranking; applied RBAC with metadata filtering for compliance-oriented multi-tenant access.
- Built modular OOP-based architecture with FastAPI REST endpoints and Streamlit frontend; containerized with Docker. Tech: Python, FastAPI, LangChain, FAISS, HuggingFace, Groq (Llama), Docker.

•Stock Analyzer Bot

RAG + LLM system for stock analysis and market trend pattern recognition.

GitHub

- Combined RAG pipelines and LLMs to analyze historical stock data and chart patterns; delivered actionable market insights via conversational AI interface.
- Tech: Python, LLMs, RAG, vector search.

•House Price Prediction API – MLOps

Production-ready ML REST API with CI/CD, Dockerized deployment, and automated testing.

GitHub

- End-to-end ML pipeline: Random Forest training, joblib serialization, FastAPI REST endpoints with Pydantic v2 validation ($R^2 \approx 0.80$); pytest unit tests automated via GitHub Actions CI/CD.
- Tech: Python, FastAPI, scikit-learn, Docker, pytest, GitHub Actions.

EDUCATION

•B.Tech in Artificial Intelligence and Data Science

Chouksey Engineering College, Bilaspur

2022 – 2026

CGPA: 8.55

CERTIFICATIONS & ACHIEVEMENTS

•AWS Solutions Architect – Associate | GeeksforGeeks

Jun 2026

•AWS Cloud Practitioner | GeeksforGeeks

Jun 2026

•4th Place – AlgoVision Data Analytics Hackathon | KOED × IIT Roorkee (national-level ML competition)

2025

•Python Certificate | HackerRank

Oct 2024

•Data Science & ML Bootcamp | Udemy

Oct 2023