

ADIBA SHAIKH

Tech Lead | Data Scientist | Gen AI & Agentic AI

adibashaikh000@gmail.com | 8879711349 | Bangalore, India

www.linkedin.com/in/adiba-shaikh-495716163/ | github.com/AdibaShaikh000

SUMMARY

Tech Lead and Data Scientist with **8 years of experience** in Python, building and shipping production GenAI systems for **Fortune 500 clients** in banking and financial services. **Led teams of 6–8 engineers** delivering enterprise AI solutions end-to-end — RAG pipelines, agentic architectures, LLM-powered chatbots, and knowledge graphs — deployed on cloud. Work spans the **full lifecycle**: problem framing with stakeholders, solution architecture, model selection and finetuning, RAG pipeline design, guardrails and PII redaction, evaluation and observability, and production deployment. The systems shipped have delivered **87% RAG accuracy** and **40–60% workflow acceleration**. Equally comfortable writing code, designing algorithms, running presales POC demos, and mentoring junior engineers.

KEY SKILLS

- **Gen AI & Agentic AI:** RAG, Hybrid Search, Reranking, AI Agents, Multi-Agent Systems, Knowledge Graph, MCP, Tool Calling, SSE Streaming, Prompt Engineering, Memory & Multi-turn
- **LLMs & Models:** GPT-4o, GPT-5.2, Llama 3, Mistral, BERT, DistilBERT, LaBSE, Transformers, Finetuning (LoRA, QLoRA)
- **Frameworks:** LangGraph, LangChain, LlamaIndex, FastMCP, HuggingFace, PyTorch, RAGAS, DeepEval
- **ML & NLP:** Classification, Random Forest, NER, BIO Tagging, Dependency Parsing, Sentence Embeddings, Semantic Reranking, Text Classification
- **Safety & Observability:** PII Detection & Redaction, Guardrails, OpenTelemetry, Azure Application Insights, LangSmith, Model Evaluation
- **Engineering:** Python, FastAPI, Flask, Django, REST APIs, Microservices Architecture, System Design, DSA, CI/CD, Docker, Git, Unit/Load/Integration Testing
- **Cloud Platforms:** Azure OpenAI, Azure DevOps, AWS Bedrock, GCP, LLM Deployment
- **Databases:** Azure SQL, SQL, MySQL, Milvus Vector DB, FAISS, Neo4j Graph DB, Azure AI Search, MongoDB, Redis, NoSQL

WORK EXPERIENCE

Synechron

05/2024 – Present | Bengaluru, India

Tech Lead

- **Led a team of 6 GenAI engineers** to architect and ship a **GPT-5.2-powered** audit chatbot on Azure OpenAI — the system auto-generates compliance test plans from regulatory documents, with **PII detection and redaction**. Integrated **OpenTelemetry with Azure Application Insights** for production observability. Compliance teams save **~18 hours/week** in manual planning
- **Led a 7-engineer team** to design and deploy a RAG chatbot on **Azure OpenAI (GPT-4o)**. Owned the full build: chunking strategy, hybrid search with Azure AI Search, semantic reranking, multi-turn conversation memory, and agentic retrieval. Shipped at **87% response accuracy** with **60% faster** document analysis
- Work **directly with stakeholders** across geographies to frame business problems, run **POC demos**, and **translate requirements into production AI solutions**.
- Architected a **Notebook-as-an-IDE framework** wired in Git version control and DAG-based workflow orchestration, cutting data science experiment cycles by **40%**

Brane Enterprises

11/2021 – 05/2024 | Hyderabad, India

Associate Solution Lead

- Built a **domain-aware knowledge graph** using **Neo4j** with **LaBSE** (Language-Agnostic BERT Sentence Embeddings) stored as native vectors, enabling intelligent entity-attribute suggestions across 20 million nodes — auto-surfacing related attributes based on solution context
- **Fine-tuned DistilBERT NER** model using **few-shot** synthetic data and BIO tagging — achieving **96% macro F1 score**; packaged as a **microservice** with **versioned REST APIs** supporting 109 languages

- Architected a **next-sentence suggestion engine** using **Milvus vector database** and **HNSW indexing**, designed as a modular, plug-in component predicting and auto-completing solution steps in real time
- **Mentored 3 junior engineers** across delivery of **8+ automation features**, establishing development standards and code review practices for the team

Tata Consultancy Services

07/2018 – 11/2021 | Mumbai, India

System Engineer

- Standardized an **enterprise document processing pipeline** using **AWS Textract and GCP Document AI** to extract structured data from multi-page PDFs, invoices, and identity documents — handling **20K+ documents monthly**
- Built **55+ production web scrapers** processing **2M+ daily records** at 95% accuracy for UK retailers
- Developed a **Random Forest classifier** to auto-categorise products into 13 categories, eliminating manual tagging.

PROJECTS

RegulationQA — Agentic AI System for US Banking Regulations

- Built a **multi-agent system** using **LangGraph and LangChain** — 4 specialized agents with ReAct loops handling Q&A on US banking regulations, with verifiable CFR citations and live data via the eCFR API. Tracked agent traces and latency with **LangSmith**
- Implemented **Graph RAG using Neo4j** to model CFR cross-references as a citation graph
- Exposed all tools as an **MCP server via FastMCP** with real-time agent streaming via **FastAPI SSE**.

RAG Assistant

- Two-stage retrieval: **FAISS bi-encoder** (top 20) + **cross-encoder reranker** (top 5) with sliding-window **conversational memory**.
- Packaged as modular **LangChain pipeline** with **LangSmith observability** and source citations in UI.

EDUCATION

Masters in Artificial Intelligence

2019 – 2021

University of Mumbai

Percentage: 86%

Bachelors in Computer Science

2015 – 2018

University of Mumbai

Grade: O (Outstanding)

CERTIFICATIONS

- **AI Engineer Agentic Track: The Complete Agent & MCP Course** - Udemy
- **AI Engineer Core Track: LLM Engineering, RAG, QLoRA, Agents** - Udemy
- **Knowledge Graphs for RAG** - DeepLearning.AI
- **ChatGPT Prompt Engineering for Developers** - DeepLearning.AI
- **AI-900: Azure AI Fundamentals** - Microsoft

AWARDS

- **Surpass Innovation Award** - Synechron
- **Special initiative Award** - TCS