FAIZAL SHAIKH

+91-98675 65991 — faizal.shaikh.eng@gmail.com — GitHub — LinkedIn — Portfolio
Nalasopara (E), Maharashtra

Summary

B.E. (CSAI & ML) graduate and National SIH 2025 Winner with hands-on experience in deep learning, LLM integration, and full-stack AI deployment. Seeking an AI Engineer role to build production-ready intelligent systems.

Education

Theem College of Engineering – University of Mumbai Boisar, Maharashtra
B.E. in Computer Science – AI & ML Specialization CGPA: 8.22/10 *Nov 2022 – May 2026*

Technical Skills

Languages & Frameworks: Python, TensorFlow, Keras, Scikit-learn, OpenCV, spaCy, FastAPI, Streamlit, React (TypeScript)
Generative AI & LLMs: Google Gemini API, Groq LLaMA 3, Prompt Engineering, LLM Integration & Evaluation, LangChain, LangGraph, FAISS, Vector Search, RAG (Retrieval-Augmented Generation), Gemini Embeddings
Data & Cloud: Pandas, NumPy, PostgreSQL, AWS, GCP, IBM Cloud, Docker, Git, Feature Engineering
Certifications: AWS – Data Engineering; NVIDIA – Disaster Risk Monitoring; IBM – AI Fundamentals & AI-Enabled Applications; Google Cloud Skills Arcade; Shell India/Edunet – Green Skills & AI; Kaggle – Intro to ML

Projects

Chronicle – Infrastructure-Grade LLMOps Platform 2025–2026
FastAPI · React (Vite) · PostgreSQL · SQLAlchemy · Groq · Tailwind CSS · Docker
- Architected immutable prompt versioning with hash-linked history and a production-promotion workflow, ensuring full auditability of LLM prompt changes.
- Engineered a Pareto Frontier optimization engine that identifies the cost-accuracy knee point for model selection; achieved 100% accuracy on sentiment benchmarks using Llama-3.3-70B.
- Built automated evaluation pipelines, an LLM-driven synthetic-data generator with Jaccard deduplication, and 180+ Pytest integration and unit tests.

SourceX – Agentic Document Intelligence & Evidence Retrieval Platform 2026
Python · Gemini API · FAISS · LangChain · LangGraph
- Built a document-grounded retrieval pipeline covering PDF extraction, text cleaning, overlapping chunking, metadata preservation, Gemini embeddings, FAISS similarity search, and query-based retrieval.
- Implemented framework-free RAG generation and a LangChain pipeline, then extended it into a LangGraph workflow with context grading, deterministic query rewriting, bounded retries, and grounded answer generation.
- Developed typed document/chunk metadata, similarity-score retrieval, automated regression tests, mocked API tests, and real Gemini embedding/retrieval smoke tests.
- Ongoing: extending the system with citation verification, evaluation, FastAPI backend, authentication, persistent storage, frontend UI, Docker, and cloud deployment.

LipSense – Deep Learning Visual Speech Recognition 2024–2025
TensorFlow · Keras · OpenCV · MediaPipe · CTC
- Developed a Conv3D encoder and bidirectional LSTM decoder trained on the GRID Corpus with CTC loss, achieving 95.2% sentence accuracy and 4.8% WER.
- Built a MediaPipe Face Mesh pipeline for real-time lip-region extraction and served inference through FastAPI.

LeafLens – AI Plant Disease Detection & Multilingual System 2024
TensorFlow · EfficientNetV2-B0 · FastAPI · Streamlit · gTTS
- Built a CNN achieving 98.4% test accuracy across 41 crop-disease classes with a 70% confidence threshold for farm-condition robustness.
- Delivered multilingual disease summaries and text-to-speech audio across 11 Indian languages with sub-two-second inference latency.

OmniText – NLP & Text Analytics Platform 2024
Python · spaCy · Streamlit · TextBlob · Sumy
- Built an NLP suite featuring extractive summarization, named-entity recognition, sentiment analysis, word clouds, and URL article extraction.

Achievements

Winner – Smart India Hackathon (SIH) 2025 Dec 2025
National Level – Government of India
- Built SimAstra, an autonomous swarm-defense framework for BEL/Central Research Laboratory, featuring multi-drone coordination and PyBullet simulation.

Member – National Service Scheme (NSS) Batch 2024–25
Theem College of Engineering