ASHNA FAISAL

AI Automation | AI Engineering | LLMs & RAG Systems

ashnahfaisal@gmail.com | +92 333 7999894 | Karachi, Pakistan | linkedin.com/in/ashna-faisal | github.com/ashna24

PROFILE

AI-focused developer specialising in workflow automation, agentic systems, and LLM-powered applications. Experienced in building n8n automations, AI agent workflows, RAG pipelines, and backend-integrated tools that connect enterprise systems through APIs, webhooks, and databases. Selected globally for the Cambridge University Summer Program (2021).

EDUCATION

BSc Computer Science *2026*

Goldsmiths, University of London (via Denning Institute, Karachi)

A-Levels & O-Levels *2017 – 2021*

The City School

TECHNICAL SKILLS

AI & Automation: n8n Workflow Automation, LangChain, LLM Agents, RAG Pipelines, Prompt Engineering, Vector Embeddings, Webhooks, REST API Integration
LLM Providers: Google Gemini API, Llama 3.2 (via Ollama), Hugging Face Transformers
Vector/Data Stores: ChromaDB, MySQL, PostgreSQL (Neon), SQL
Languages: Python, JavaScript, SQL
Tools & Frameworks: FastAPI, Docker, Git, NumPy, pytest, PaddleOCR

EXPERIENCE

AI Engineer Intern, Al Rahim Group of Industries *1 October 2026 – 10 November 2026*

- Built an n8n workflow with PostgreSQL and a Gemini AI agent for the company's project management dashboard that extracts requirements from project PDFs, generates tasks, and assigns them by employee skills and availability, with a Gmail human-in-the-loop step for managers to approve or reject tasks before dashboard submission.
- Automated daily update reminders, weekly project reports, and deadline monitoring for the project management dashboard using n8n and PostgreSQL.
- Built a WhatsApp AI assistant using Python, FastAPI, the Gemini API, and Neon PostgreSQL, enabling natural-language queries for employee tasks and daily updates.
- Integrated the assistant with the task-generation workflow through PDF intake, so project documents sent on WhatsApp trigger automated task creation.

Vice President, Denning Beyond Tech Society *2023*

- Co-founded Denning Institute's first tech society and orchestrated a 100+ participant coding competition, establishing a recurring annual event.

Freelance Web Developer, VisaBridge Global (Visa/Immigration Consultancy) *2026*

- Designed and launched a live business website (visabridge.live) for a visa/immigration consultancy, establishing its first online presence using Google AI Studio.

PROJECTS

AI Customer Support & Lead Automation Pipeline

Stack: n8n, Google Gemini API (gemini-2.5-flash-lite), Google Sheets API, Gmail API, JavaScript, Docker, Webhooks

- Built a production-ready AI customer support automation workflow in n8n that reduced per-ticket handling time from approximately 2 minutes to under 5 seconds by fully automating message classification, priority routing, sentiment detection, and personalised reply generation.
- Integrated Google Gemini AI via n8n's AI Agent nodes, webhook input validation, JavaScript JSON parsing with error recovery, and real-time Google Sheets audit logging in a self-hosted Docker environment.

Lex — AI-Powered Legal Document Analyser *Final Year Project*

Stack: Python, Ollama, Llama 3.2, Hugging Face Transformers, PaddleOCR, pytest, Git

- Built an end-to-end pipeline that extracts text from scanned legal documents via OCR, splits it into individual clauses using rule-based NLP, classifies each clause's risk level using Llama 3.2, and optionally translates plain-English explanations into Urdu.
- Evaluated and selected models through structured comparison testing at each pipeline stage — OCR (PaddleOCR vs. EasyOCR), classification (Llama vs. BART/DeBERTa), and translation (NLLB-200 vs. Helsinki-NLP vs. Llama) — including native-speaker verification of translation accuracy.
- Validated pipeline reliability with 36+ automated tests (pytest) and version-controlled, iterative development via Git.

RAG-Based Academic Course Assistant

Stack: Python, Streamlit, LLMs, Google Gemini (gemini-1.5-flash), gemini-embedding-001, NLP, pypdf, NumPy, pickle

- Built an end-to-end RAG pipeline using gemini-embedding-001 vectors and NumPy cosine similarity search, enabling semantic Q&A over course PDFs with no external vector database.
- Cut redundant embedding API calls to zero by caching vectors to disk with pickle and detecting existing embeddings on startup.
- Added a similarity score threshold to automatically detect and reroute low-confidence queries to a fallback response before passing them to the language model, reducing hallucinated answers.