

ABDULLAH USMAN

Software Engineer | AI / ML Engineer | Generative AI Developer | Data Science
+92 334 027 4581 | abdullahu2003@gmail.com | [LinkedIn](#) | [Github](#)

PROFESSIONAL SUMMARY

Results-driven Software Engineering graduate with hands-on experience delivering production-grade AI systems across enterprise and startup environments. Specialises in LLM integration, Retrieval-Augmented Generation (RAG), NLP pipelines, and computer-vision solutions at scale. Comfortable taking projects from raw data to deployed application — combining strong Python fundamentals with modern GenAI tooling, REST API design, and cloud deployment. Seeking a junior / associate AI or ML engineering role where deep technical ownership and impact-driven development are valued.

EXPERIENCE

AI / ML Intern — Sui Southern Gas Company (SSGC)

DEC 2025] – FEB 2026] | Karachi, Pakistan

- Designed and deployed an end-to-end AI image-processing web application handling **3.2M+ enterprise images**, replacing a fully manual classification workflow.
- Built a scalable ML pipeline to automatically classify, sort, and organise images into structured company directories — reducing manual intervention significantly.
- Implemented OCR-based information extraction with open-source libraries to digitise sensitive internal documents securely.
- Ensured 100 % compliance with internal data-security policies using exclusively open-source technologies.
- **Stack:** Python, Streamlit, OpenCV, EasyOCR, Tesseract, Scikit-learn, FastAPI, Plotly.

AI / ML Intern — ITSOLERA

FEB 2026 – APRIL 2026] | Remote

- Collaborated in a team of 4 engineers to build a production NLP web application for real-time customer sentiment classification.
- Processed and cleaned **1M+ customer reviews**; performed EDA and feature engineering to surface actionable sentiment patterns.
- Trained and optimised ML classifiers (TF-IDF + Scikit-learn) achieving strong categorisation accuracy across multi-class sentiment labels.
- Integrated trained model into a Django web app with real-time inference endpoint and an interactive KPI analytics dashboard.
- **Stack:** Python, Scikit-learn, Pandas, Matplotlib / Seaborn, Streamlit, FastAPI.

EDUCATION

NED University of Engineering & Technology, Karachi

Bachelor of Software Engineering | 2021 – 2025

CERTIFICATIONS

- Oracle OCI AI Foundations Associate
- Oracle OCI Generative AI Professional
- Hugging Face AI Agent Certification

PROJECTS

AI-Powered Inventory Management System

Final Year Project — NED University | Python · Django · LSTM · XGBoost · React Native · Next.js · MongoDB · DigitalOcean

- Built a full-stack AI inventory platform with demand forecasting and automated stock replenishment for retail use cases.
- Trained LSTM, XGBoost, and regression models on historical sales data; evaluated with RMSE & MAE for iterative performance optimisation.
- Implemented low-stock detection and intelligent auto-reorder logic with REST API backend serving both React Native (mobile) and Next.js (web) frontends.
- Deployed to DigitalOcean with CI/CD pipeline for automated testing and release.

Local LLM Chatbot — Privacy-First RAG Assistant

Personal Project | Ollama · LLaMA 3 · LangChain · ChromaDB · Python

- Developed a fully offline AI assistant using LLaMA 3 via Ollama with persistent memory and contextual multi-turn responses.
- Implemented vector embeddings + ChromaDB retrieval for improved context handling across long conversations.
- Designed modular RAG-ready architecture allowing easy swap of embedding model or vector store.

AI Pizza Review Assistant (RAG System)

Personal Project | LangChain · ChromaDB · Vector Embeddings · Python

- Built a Retrieval-Augmented Generation pipeline using embeddings and similarity search for domain-specific QA.
- Optimised chunking strategy and retrieval pipeline to improve answer relevance and reduce hallucination.

Generative AI Healthcare Translator

Personal Project | Python · LLM APIs · Speech Recognition · Translation APIs

- Built a multilingual AI system enabling real-time doctor–patient communication across low-resource languages.
- Integrated speech-to-text, automatic language detection, and LLM-based translation refinement for context-aware responses.

NLP Customer Review Classification System

ITSOLERA | Python · Scikit-learn · TF-IDF · Django · Pandas

- Processed 1M+ reviews with full NLP preprocessing pipeline (tokenisation, lemmatisation, stop-word removal) and TF-IDF feature engineering.
- Deployed model to Django web app with real-time inference and analytics dashboard.

TECHNICAL SKILLS

Languages:	Python (primary), C#, SQL
Frameworks & APIs:	Django, FastAPI, Flask, Streamlit, REST APIs
AI / ML & Deep Learning:	Scikit-learn, PyTorch, LSTM, XGBoost, Supervised Learning, Classification, Regression, Time-Series Forecasting, Feature Engineering, Hyperparameter Tuning
Generative AI & LLM Engineering:	LangChain, Ollama (LLaMA 3), RAG, Vector Embeddings, Prompt Engineering, Context Management, LLM API Integration
NLP:	Text Preprocessing, TF-IDF, Tokenization, Lemmatization, Sentiment Analysis, Text Classification
Computer Vision & OCR:	OpenCV, Tesseract OCR, EasyOCR, Image Classification, Batch Processing Pipelines
Vector Databases & Storage:	ChromaDB, FAISS, Supabase, SQL, MongoDB
Data Science & Analytics:	Pandas, NumPy, EDA, Data Cleaning, Matplotlib, Seaborn, Plotly, Forecasting, KPI Dashboards
Cloud, DevOps & Deployment:	Git, GitHub, CI/CD Pipelines, Docker (basics), DigitalOcean, Vercel, Streamlit Cloud, AWS (basics)
Automation & Workflow:	n8n, REST API Integration, AI Workflow Orchestration