

Muhammad Nouman Hanif

AI ENGINEER | DATA SCIENTIST | CLOUD SPECIALIST

Faisal Town, Lahore, Pakistan

☎ +92 348 2203515 | ✉ mnoumanhanif66@gmail.com | 📷 mnoumanhanif | 🌐 mnoumanhanif

Summary

Data Science and **AI Engineer** with expertise in **Computer Vision** (ViT, Swin, YOLO), **Generative AI** (LLMs, QLoRA, PEFT), and end-to-end **machine learning** pipelines. Experienced in **fine-tuning 7B+ parameter models**, building **RAG** and **LangChain** systems, and developing evaluation frameworks for medical AI. Strong background in **MLOps** with AWS, Docker, CI/CD, and cloud-scale deployment. **Gold Medalist Electrical Engineer** with 6+ years of **DevOps** and automation experience, including 100+ **RPA** bots and production-grade AI implementations.

Education

Master of Science in Data Science (MS-DS)

Lahore, Pakistan

NATIONAL UNIVERSITY OF COMPUTER & EMERGING SCIENCES (FAST-NUCES)

Jan. 2024 – Present

- **CGPA:** 3.83 / 4.0
- **Specialization:** Data Science, Deep Learning, Computer Vision, & Generative AI.
- **Thesis Research:** Advanced Research in Medical GenAI for Medical Vision-Language Models using PyTorch and Transformers.

Bachelor of Science in Electrical Engineering (BS-EE)

Karachi, Pakistan

NATIONAL UNIVERSITY OF COMPUTER & EMERGING SCIENCES (FAST-NUCES)

Aug. 2019 – Jul. 2023

- **CGPA:** 3.86 / 4.0 (Gold Medalist)
- **Specialization:** Computer Engineering, Embedded Systems, & Power Systems.
- **FYP Project:** IoT-Enabled Smart Solar Inverter with MPPT (Leader). Integrated embedded C++ control systems with real-time monitoring.

Professional Experience

National University of Computer & Emerging Sciences (FAST-NUCES)

Lahore, Pakistan

INSTRUCTOR - COMPUTER SCIENCE (DS, CYS, AI)

Aug. 2023 – Present

- Taught core computing courses: Programming Fundamentals, Data Structures, OOP, DLD, COAL, and Computer Networks.
- Developed interactive course materials, assignments, and assessments to improve student engagement.
- Mentored students in algorithm design, coding best practices, and technical problem-solving.

syslab.ai

Karachi, Pakistan

CLOUD ENGINEER INTERN

June 2024 – Aug. 2024

- Migrated applications (MongoDB, NGINX, Mail) to AWS for improved performance and scalability.
- Implemented security measures including Cloudflare Web Application Firewall (WAF), SSL/TLS encryption, and AWS IAM policies.
- Automated server management and CI/CD pipelines using AWS CloudWatch and infrastructure as code (IaC) principles.

Woot Tech Aerospace

Karachi, Pakistan

ENGINEERING INTERN

Mar. 2023 – Apr. 2023

- Contributed to drone manufacturing and aircraft structural design projects.
- Designed 3D models and technical schematics using SolidWorks and AutoCAD.

National Transmission & Despatch Company (NTDC)

Hyderabad, Pakistan

ELECTRICAL ENGINEERING INTERN

Jul. 2022 – Sep. 2022

- Gained hands-on experience in power distribution, grid operations, and electrical safety procedures.

Artistic Milliners

Karachi, Pakistan

ENGINEERING INTERN

Jun. 2022 – Jul. 2022

- Assisted in maintenance of industrial power systems, including generators, gas turbines, and boilers.
- Learned about large-scale power management in a manufacturing environment.

Technical Skills

Deep Learning & CV	PyTorch, TensorFlow, Vision Transformers (ViT, Swin), Stable Diffusion YOLO, OpenCV, CNNs, GANs, VAEs, CLIP, Grad-CAM, Attention Rollout
Generative AI & LLMs	LLM Fine-tuning (LoRA, QLoRA, IA3), Hugging Face, RAG, LangGraph (Multi-Agent), LLM-as-a-Judge, LLaMA 3, Mistral 7B, T5
MLOps & Infrastructure	AWS, Azure, GCP, DigitalOcean, Docker, Kubernetes, CI/CD Pipelines, NGINX, Linux/Bash, Git, GPU Optimization (4-bit quantization, PagedAdamW)
Languages & Data	Python, C++, SQL, PySpark, SQL, MongoDB, Data Curation, Pandas, NumPy, Scikit-Learn

Projects & Research

Secure Face Anti-Spoofing System with Vision Transformers

PyTorch, Swin, ViT

COMPUTER VISION SECURITY PROJECT

2025

- Developed a robust biometric security system to detect face spoofing (replay/print attacks) using Transformer architectures.
- Fine-tuned **Swin Transformer** achieving **99.75% test accuracy** and near-perfect recall on the CelebA-Spoof dataset.
- Benchmarked against ViT (99.37%), proving Swin's hierarchical attention superiority for detecting subtle spoofing artifacts.

Brain Tumor Classification with Explainable AI (XAI)

ViT, Grad-CAM

MEDICAL IMAGING AI PROJECT

2025

- Engineered a multi-class classification model (Glioma, Meningioma, Pituitary) achieving **99.65% accuracy** using Vision Transformers.
- Implemented **Grad-CAM** and **Attention Rollout** to visualize decision heatmaps, verifying the model focused on pathological regions rather than skull artifacts, ensuring clinical reliability.

Parameter-Efficient Fine-Tuning (PEFT) & Agentic Systems

QLoRA, LangGraph, LLaMA-3

LLM RESEARCH & OPTIMIZATION

2025

- Conducted a 42-run ablation study on DeBERTa-v3, proving **QLoRA** reduced VRAM usage by **~90%** (7.2GB → 1.5GB) while retaining **97%** of Full-FT accuracy on Sentiment Analysis tasks.
- Identified that increasing Learning Rate to 2×10^{-5} and targeting Feed-Forward Network (FFN) layers were critical for PEFT convergence, outperforming rank inflation ($r = 16$).
- Built a **Multi-Agent Research Assistant** using **LangGraph**, orchestrating specialized agents (Keyword, Search, Rank, Summarize) to automate literature review, evaluated using ROUGE and Factuality scores.

Multimodal Pipeline: Story Generation & E-Commerce Mockups

Stable Diffusion, BLIP, T5, LoRA

GENERATIVE VISION SYSTEM

Jun. 2025

- Engineered an end-to-end multimodal pipeline combining **BLIP** (Image Captioning) and **FLAN-T5** (Story Generation), fine-tuned via LoRA on the Flickr30k dataset to generate creative narratives from visual inputs.
- Built an automated product mockup system using **Stable Diffusion (img2img)**, optimizing 'denoising_strength' (0.85) and 'guidance_scale' to overlay photorealistic textures (e.g., floral, Art Deco) onto product geometry while maintaining structural fidelity.

Real-Time Object Detection & Generative Model Analysis

YOLOv8, GANs, VAEs, OpenCV

DEEP LEARNING SYSTEMS

2024

- Fine-tuned **YOLOv8s** on a custom COCO subset, managing the full lifecycle from JSON-to-TXT formatting to training (10 epochs) and mAP50 validation. Deployed in a real-time OpenCV inference loop.
- Benchmarked Generative models on MNIST/CIFAR-10: Proved **GANs** achieved superior image fidelity (FID: 24.26) over VAEs. Repurposed VAE for Anomaly Detection on Fraud data, achieving **0.93 F1-Score**.

Flood Severity Prediction with Ensemble Gradient Boosting

CatBoost, XGBoost, Scikit-Learn

PREDICTIVE MODELING & MACHINE LEARNING

Dec. 2024

- Engineered a robust flood forecasting pipeline using **CatBoost** and **XGBoost**, trained on a massive dataset of **1.1 million+ records** encompassing 22 meteorological and topographical features.
- Achieved a high predictive accuracy of **0.85 R^2** , conducting rigorous feature importance analysis to identify Rainfall and Soil Moisture as critical predictors for real-time early warning systems.

Transformer-Based Biological Sequence Classification

PyTorch, BERT, Hugging Face

DEEP LEARNING RESEARCH

Jun. 2024

- Implemented a Transformer-based deep learning architecture to classify complex biological sequences, leveraging **BERT** embeddings for feature extraction.
- Fine-tuned the model using **PyTorch** and Hugging Face 'transformers', optimizing hyperparameters to maximize classification accuracy on genomic data.

High-Scale RPA Automation Suite

Python, UiPath

AUTOMATION ENGINEERING

2019 - 2023

- Architected and deployed **100+ automation bots** to handle complex workflows (YouTube engagement, file management), reducing manual workload by **> 80%** through headless browser automation.
- Integrated Computer Vision-based selectors for mobile testing bots, enabling robust UI interaction for dynamic applications.

IoT-Enabled Smart Solar Inverter with MPPT

C++, Arduino, IoT

EMBEDDED SYSTEMS ENGINEERING (FYP)

Aug. 2022 - Jul. 2023

- Lead Engineer for a Smart PV Inverter system. Developed **C++ firmware** for real-time Maximum Power Point Tracking (MPPT), optimizing energy conversion efficiency.
- Integrated hardware sensors with an ESP32 microcontroller to transmit real-time telemetry to a cloud dashboard, demonstrating full-stack IoT capability.

Honors & Awards

2019 - 2023 **Gold Medalist**, National University of Computer and Emerging Sciences (NUCES-FAST)

Karachi, Pakistan

2022 - 2023 **Appreciation Award**, serving as the Director of the Hardware Exhibition at PROCOM

Karachi, Pakistan

2019 - 2023 **Dean's & Rector's List**, National University of Computer and Emerging Sciences (NUCES-FAST)

Karachi, Pakistan

2015 - 2017 **1st Position**, New Cantt Public High School

Lahore, Pakistan

Certifications

Generative AI & LLMs

Oct. 2025 **Generative AI Engineering & Fine-Tuning Transformers**, IBM - Coursera

Aug. 2025 **Fundamentals of AI Agents Using RAG and LangChain**, IBM - Coursera

Aug. 2025 **Gen AI Foundational Models for NLP & Language Understanding**, IBM - Coursera

Mar. 2023 **Crash Course on Python**, Google - Coursera

DevOps & Automation

Sep. 2023 **Introduction to DevOps**, IBM - Coursera

Jul. 2023 **Introduction to Cloud Computing**, IBM - Coursera

Jan. 2022 **Robotic Process Automation (RPA) Basics and UiPath**, UiPath - Coursera

Professional Certifications

Oct. 2023 **Registered Electrical Engineer**, Pakistan Engineering Council

Oct. 2023 **PEC Generative AI Training**, Pakistan Engineering Council