

# Ismail Qayyum

+92 334-020-2266 | ✉ ismail.qayyum4@gmail.com | [in ismailqayyum](#) | [G IsmailQayyum](#) | [M @ismail.qayyum4](#) | [📍 Karachi, Pakistan](#)

## SUMMARY

AI Engineer with a B.S. in Artificial Intelligence (FAST-NUCES) and production-proven experience designing end-to-end data pipelines, LLM-powered agentic systems, and ML-driven automation across cloud-native environments. Hands-on with Apache Airflow, Spark, Medallion Architecture, OpenAI API, Claude API, LangChain, and LangGraph. Authored a published computer vision dataset (Zenodo DOI) and technical articles in *AI Monks*. Delivered measurable impact: 45% faster queries, 80% reduced deployment time, 60% less manual effort, 30% accuracy improvement.

## SKILLS

**Data Engineering** : Apache Airflow, Apache Spark, Medallion Architecture (Bronze/Silver/Gold), AWS S3, MongoDB, Neo4j, OpenSearch, Indexed Views, ETL/ELT Pipelines

**Generative AI** : LangChain, LangGraph, LlamaIndex, RAG Pipelines, Agentic Orchestration, OpenAI API, Claude API, Prompt Engineering, LLM Fine-tuning, Hugging Face

**Deep Learning** : PyTorch, TensorFlow, CNNs, RNNs/LSTMs, GANs, VAEs, Diffusion Models, Attention Mechanisms

**ML & Data Science** : Scikit-learn, NumPy, Pandas, Matplotlib, NLTK, NLP, Computer Vision, YOLO, OpenCV

**Languages** : Python, Java, C++, Dart

**Cloud & Platforms** : AWS, Microsoft Azure, PowerApps, Hugging Face Hub

**DevOps & Tools** : Docker, Kubernetes, FastAPI, Redis, Git, n8n, Make (Integromat), Airtable, REST API Design, OpenAPI

## EXPERIENCE

### Disrupt.com (formerly Gaditek)

Jun. 2025 – Present

Karachi, Pakistan

#### Software Engineer II — AI & Data Systems

Apr. 2026 – Present

- Architected end-to-end **Medallion Architecture pipelines** (S3 → Bronze → Silver → Gold) using **Apache Airflow** and **Apache Spark**, delivering features from requirements to production across optimised **MongoDB indexed views** and REST APIs
- Owned full feature lifecycle, from data ingestion through pipeline orchestration, schema design, API delivery, and product shipping — across Secure.com's **10,000+** AWS asset security platform

#### Software Engineer I — AI & Data Systems

Jun. 2025 – Apr. 2026

- Engineered real-time ingestion pipelines processing **1M+ daily events per tenant** from AWS, OpenSearch, and Neo4j, powering multi-tenant security analytics for Secure.com
- Integrated LLM based automation for intelligent error detection and self-correction across 300K+ vulnerability records, improving data accuracy by **30%**
- Designed and optimised multi-environment databases managing 300K+ CVE/CWE entries; achieved **45% query performance gain** across staging and production
- Developed data sync tooling to automate staging → production movement, cutting deployment time by **80%** and manual review workload by **60%**

### Imperium Dynamics

Apr. 2025 – Jun. 2025

Machine Learning Engineer Intern

Karachi, Pakistan

- Designed and shipped LangChain + FastAPI **autonomous agents** for enterprise automation with chained LangGraph multi-step reasoning logic deployed in live client environments
- Implemented a **Redis-backed stateful memory layer** for low-latency, high-availability LLM interactions in production, achieving infrastructure-grade reliability
- Automated internal email workflows via Microsoft PowerApps, reducing manual operations by **40%**
- Authored FastAPI service architecture following enterprise standards: security, modularity, versioning, and full OpenAPI documentation

### FAST-NUCES

Jan. 2023 – Apr. 2023

Student Research Lead — Computer Vision & Generative AI

Karachi, Pakistan

- Led creation of **NBF\_StreetView** — a 19,000 image-pair dataset for autonomous vehicle sensor simulation; **published on Zenodo (DOI: 10.5281/zenodo.14552482)**
- Extended a diffusion model with stage-wise diffusion, adaptive conditioning, and attention mechanisms for real-world image degradation restoration
- Designed a novel composite loss function (Contextual + Perceptual + L1) that measurably improved reconstruction fidelity and realism

## PROJECTS

---

**Teaching Assistant Agent** | [LangChain](#) · [LangGraph](#) · [FastAPI](#) · [LLM Agents](#) [GitHub](#)

- Autonomous LLM agent that grades student assignments end-to-end — parses PDFs, evaluates answers against rubrics, and generates structured feedback reports; eliminates hours of manual grading per session

**NBF\_StreetView Dataset** | [Python](#) · [OpenCV](#) · [Imageaug](#) · [Zenodo](#) [GitHub](#) | [DOI](#)

- Published computer vision dataset of 19,000 degraded/clean street-view image pairs for autonomous vehicle sensor research — cited as Qayyum, I., Badar, A., & Rizwan, A. (2024), Zenodo

## PUBLICATIONS & WRITING

---

**NBF\_StreetView Dataset** | [Zenodo](#) · [DOI: 10.5281/zenodo.14552482](#)

- Qayyum, I., Badar, A., & Rizwan, A. (2024). *NBF\_StreetView* [Data set]. Zenodo. Published computer vision research dataset of 19,000 degraded/clean street-view image pairs

**Autoencoders: Basics and Beyond** | *AI Monks* · [Jun 2024](#) [medium.com/@ismail.qayyum4](#)

- PyTorch implementation of autoencoders from scratch; published in *AI Monks* publication on Medium

**Understanding The Neural Network Model** | *AI Monks* · [Oct 2023](#) [medium.com/@ismail.qayyum4](#)

- Deep dive into neural network architectures and forward propagation; published in *AI Monks* on Medium

## EDUCATION

---

**National University of Computer and Emerging Sciences (FAST-NUCES)**

Sep. 2021 – Jun. 2025

*Bachelor of Science in Artificial Intelligence*

Karachi, Pakistan

- **Relevant Coursework:** Machine Learning, Deep Learning, Computer Vision, NLP, Data Structures & Algorithms, Statistics, Database Systems
- **Research:** NBF\_StreetView dataset (published author, Zenodo 2024); Diffusion model research for image restoration