

# Talal Muzaffar

LinkedIn: [linkedin.com/in/talalmuzaffar](https://www.linkedin.com/in/talalmuzaffar)

GitHub: [github.com/talalmuzaffar](https://github.com/talalmuzaffar)

Portfolio: [talal-portfolio-seven.vercel.app](https://talal-portfolio-seven.vercel.app)

Email: [talalmuzaffar3@gmail.com](mailto:talalmuzaffar3@gmail.com)

Mobile: +92-304-7919692

## EDUCATION

---

- **FAST National University of Computer and Emerging Sciences** Islamabad, Pakistan  
*Bachelor of Science - Data Science* 2020 – 2025

## TECHNICAL PROFICIENCIES

---

- **Generative AI:** PyTorch, Large Language Models (LLaMA, Falcon, Gemma, Phi, BART), LangChain, LlamaIndex, Fine-tuning (LoRA, QLoRA)
- **Agentic AI:** Autonomous Agent Workflows, RAG, Text-to-SQL, LangGraph, Google ADK, MCP
- **Data Extraction:** LayoutLM, DONUT, Azure Document Intelligence, PaddleOCR, NER Pipelines
- **Data & Infra:** Microsoft Azure, GCP, MongoDB, ChromaDB, Qdrant, MySQL
- **MLOps:** FastAPI, Docker, GitHub Actions, Jenkins, Airflow, MLFlow

## EXPERIENCE

---

- **DOST AI** Barcelona, Spain (Remote)  
*Machine Learning Lead* January 2026 – Present
  - Engineered an autonomous onboarding agent that fetches organization metadata, parses customer-uploaded CSVs, and auto-maps vendor field names — across **100+ vendors per customer with 50–1000+ fields each** — to the internal database schema, cutting mapping time **10 min → 5 min** per batch with a human-in-the-loop final review gate.
  - Scaled the pipeline to onboard **10+ enterprise customers per month**, eliminating a previously manual, per-vendor field-reconciliation process that did not scale beyond dozens of mappings.  
*Machine Learning Engineer* March 2024 – December 2025
  - Architected a production RAG pipeline over a **3.8M+ document** corpus of invoices, purchase orders, and delivery notes, enabling natural language financial queries — vendor dependency analysis, invoice due-date forecasting, monthly spend projections — with **sub-second retrieval latency**.
  - Built and deployed FastAPI microservices processing **40K documents / 100K+ pages per month**, performing full-field entity extraction — vendor names, IDs, PO numbers, due dates, payment terms, line-item details, taxes, and totals — feeding a structured financial knowledge base.
- **KDD Lab, FAST NU** Islamabad, Pakistan  
*Machine Learning Engineer* June 2023 – December 2023
  - Designed an LLM-powered resume validation system that scrapes applicant GitHub profiles, analyzes repository languages and architecture via LLM reasoning, and automatically flags skill mismatches between claimed and demonstrated competencies — built as a core module of a larger AI hiring engine.
  - Fine-tuned domain-specific LLMs (LLaMA, Falcon, BART, Phi, Gemma) on medical corpora to automate specialized enterprise workflows.
- **FAST National University of Computer and Emerging Sciences** Islamabad, Pakistan  
*Teaching Assistant* January 2023 – April 2023
  - Fundamentals of Software Engineering
  - Information Security

## PROJECTS

---

- **UniversityTHRIVE | Freelance Developer** November 2025 – January 2026
  - Built a KNN-based researcher matching platform — deployed at **UNCW and multiple US universities** — automating group formation by career stage, discipline, department, and availability, eliminating manual facilitator assignment and Zoom scheduling overhead.

- **Technologies:** Next.js, FastAPI, KNN Matching, SMTP
- **Medical Abstract Generation**
  - Fine-tuned **Falcon 7B** with **QLoRA** on a T4 GPU using MICCAI proceedings to generate structured abstracts from full research manuscripts, achieving a **BERT Score of 0.90**.
  - Applied 4-bit quantization to fit a 7B-parameter model within 16GB VRAM constraints while preserving generation quality.
  - **Technologies:** PyTorch, HuggingFace Transformers, QLoRA, Falcon 7B
- **Student Counsellor Chatbot**
  - Built a conversational RAG agent over a university academic knowledge base to answer student queries on resources, policies, and procedures.
  - **Technologies:** LangChain, LLMs, RAG

## PUBLICATIONS & HONORS

---

- **Publication:** Choices for Measuring Cognitive Function Digitally in Low- and Middle-Income Countries. *American Epilepsy Society*, 2023.
- **Dean's List:** Placed on the Dean's List, FAST National University (2024–2025).
- **McKinsey Forward Program:** Selected participant (McKinsey & Company, 2021).
- **Crossroads Emerging Leaders:** Selected participant (Harvard University, 2020).
- **Intelligent Homes & Health Challenge:** Finalist (The New York Academy of Sciences, 2019).