# MUHAMMAD ABDULLAH HAROON

+92 306 0590604 | abdullaharoon.work@gmail.com | LinkedIn | GitHub | Google Scholar | Portfolio

## EDUCATION

**National University of Computer and Emerging Sciences (FAST-NUCES)** — Lahore, Pakistan
Bachelor of Science in Computer Science — Aug 2023 – Expected 2027

## EXPERIENCE

**FAST School of Management** — Sep 2026 – Present
Research Assistant, AI in Healthcare Management — Lahore, Pakistan

- Surveying **50–60 doctors**, then their patients, on how they use **AI** at work and at home, and whether they see it as an opportunity or a threat. Supervised by Dr. Beenish Arshad.
- Leading to the analysis of data in **Python dashboards** to scope and development of an **AI healthcare management system** for scheduling, resource allocation, OPD Inventory Management, and patient–doctor assignment, proposed for HEC NRPU.

**FAST School of Computing** — Jul 2026 – Sep 2026
Research Intern, supervised by Dr. Maryam Bashir — Lahore, Pakistan

- Showed that Urdu fake-news models drop from **0.771 to 0.005 macro-F1** across datasets because they learn article length, then found the same shortcut in English fake news and sarcasm.
- Built a **26-dataset, 188,309-example** Urdu benchmark in Nastaliq and Roman Urdu with cluster-aware splits so near-duplicates never leak between train and test.
- Across **80** domain shifts, fine-tuned encoders lost **13.5 – 47.6 macro-F1**, while few-shot **LLMs** won **30 of 42** comparisons under Holm-corrected bootstrap tests over 3 seeds.

**QuickFuel** — Jun 2026 – Jul 2026
IT Intern — Lahore, Pakistan

- Introduced the company's first **AI/ML** work; auditing the .NET/WPF Management Portal put **9 AI use cases** and **6 design flaws** before senior leadership.
- Prototyped **LLM** sentiment analysis (**FastAPI, PostgreSQL**) and a **ChromaDB** search that finds related products, not just keyword matches.

**Optimus Automate** — Jun 2026 – Jul 2026
AI Automation Intern — Remote

- Built an **LLM agent** on the **Gmail API** that classifies mail by intent and drafts replies.
- Built a **RAG** tool that answers and summarizes PDFs, citing only text it actually retrieved.

## RESEARCH

**Domain Robustness of Multilingual NLP Models Across Urdu and Roman Urdu Scripts**
M. A. Haroon, M. Bashir | *In preparation, Springer Language Resources and Evaluation (Q1)* | GitHub

- Released **26** datasets (**188,309** examples) as **7** transfer matrices over **4** task types, each documented with origin and licence (Zenodo DOI).
- A target-trained model wins **83–100%** of shifts, so encoder losses are robustness failures; the worst come from domains labelling different text the same way.

**Does Compression Preserve Adversarial Patch Robustness? An Empirical and Certified Study of Pruned, Quantized, and Distilled CNNs**
M. A. Haroon, A. R. Antall, A. Hadi | *In preparation, TMLR* | GitHub

- First side-by-side test of **pruning, quantization and distillation** against patch attacks: **13** ResNet-18 variants at **2**/**4**/**8**×, with a dense-small control to separate sparsity from capacity.
- Covers universal and **LaVAN** patches, a **13**×**13** transfer matrix and the first **PatchCleanser** certified accuracy on compressed models, over 3 seeds.

**When Benchmarks Mislead: Shortcut Learning, Length Confounds, and the Limits of Cross-Dataset Generalization in Multilingual Fake News and Sarcasm Detection**
M. A. Haroon | arXiv:2607.14131 | GitHub

- Fake Urdu articles run **3.4**× longer than real ones, sinking transfer to **0.005 macro-F1**; the pattern repeats on English fake news (WELFake vs. ISOT) and sarcasm (TweetEval vs. Sarcasm Corpus V2).

**Bitcoin Price Direction Prediction via Regime-Aware Multi-Modal Fusion of Social Sentiment and Technical Features**
M. A. Haroon | arXiv:2607.23370 | GitHub

- **RAML** fuses Reddit sentiment (FinBERT) with price data (BiLSTM) conditioned on 24-hour volatility, reaching **0.547** / **0.551 macro-F1** at 3h / 6h on **3,491** hourly observations.

## PROJECTS

**Sachai Check (Final Year Project)** | Python, FastAPI, React, Hugging Face, PostgreSQL *Ongoing*

- Detects Urdu **fake news** across domains and **deepfake audio** in two pipelines; an LLM transliterates Roman Urdu first.
- An **LLM judge** backs each verdict with evidence and a confidence score instead of a bare real/fake label.

**TaxChronos** | LangChain, Qdrant, Gemini API, RAGAS GitHub

- Dense-retrieval **RAG** for mixed English–Urdu tax questions: **under 20s** vs. 2–3 min by hand, scored on **RAGAS** faithfulness and context recall.

**Bilingual WhisperBench** | Python, OpenAI Whisper, FastAPI GitHub

- Tested Whisper Tiny, Base and Small on **327** Urdu–English clips from **11** speakers; new **SPER** metric; best **WER 0.899**.

**CRMind** | React 19, Node.js, Express, MongoDB, Gemini API GitHub

- Sales CRM with **30 REST endpoints**, **JWT** auth, a **5-stage** drag-and-drop pipeline and a **4-KPI** dashboard.
- **Gemini** rates lead risk **0–100**, drafts emails and reviews pipeline health, with a **30s timeout** and fallback model.

**Daily Digest** | FastAPI, PostgreSQL, Gemini API, Docker GitHub

- Live news digest for **15+ active users**: fetches, ranks by Pakistan relevance, summarizes with **Gemini** and emails it with no manual step.

## TECHNICAL SKILLS

**Languages:** Python, C, C++, C#, SQL, JavaScript, TypeScript, HTML, CSS

**ML & GenAI:** PyTorch, TensorFlow, Scikit-learn, Hugging Face, LLMs, LangChain, RAG, AI Agents, NLP, Whisper

**Data & Databases:** Pandas, NumPy, Matplotlib, PostgreSQL, MS SQL Server, SQLite, MongoDB, Qdrant, ChromaDB, FAISS

**Web & Tooling:** FastAPI, Node.js, Express.js, React.js, REST APIs, SQLAlchemy, .NET, Git, Docker, Jupyter, n8n

## CERTIFICATIONS

**AI Engineer Core Track: LLM Engineering, RAG, QLoRA, Agents** – Udemy Certificate

**Machine Learning with Python** – IBM Certificate

**Data Analysis and Visualization with Python** – Microsoft Certificate

**Introduction to Generative AI** – Google Certificate