

HAMZA MOORAJ

AI Researcher — Multimodal & Agentic Systems | Evaluation, Reliability, and Generalisation

Pakistan | +92 317 3499268 | hhmooraj@gmail.com | [Portfolio](#) | [GitHub](#) | [LinkedIn](#) | [HuggingFace](#) |

PUBLICATIONS

A Systematic Exploration of Architectural Trade-offs for Crop Disease Classification - Lead Author

July 2025 – April 2026 — Accepted to Transactions on Machine Learning Research (TLMR), 2026 — [arXiv Preprint](#)

- Created AgriPath-LF16, a 65-class benchmark with explicit lab/field domain separation.
- Designed a unified evaluation framework comparing CNNs, contrastive VLMs, and generative VLMs.
- Analysed model behaviour under domain shift, evaluating robustness across architectures.
- Provides empirical evidence on model behaviour under domain shift across CNN and VLM architectures.

PROJECTS

Mechanistic Analysis of Induction in a Tiny Transformer - Personal Project

February 2026

- Trained a 2-layer transformer on a synthetic induction task to study learned behaviour.
- Analysed attention patterns to characterise offset-diagonal and recency-biased mechanisms.
- Designed interventions (ablation, rotation, perturbation) to probe causal structure.
- Identified shortcut learning behaviour in place of canonical induction mechanisms.

Medical Question-Answering (QA) Conversational Agent (Course: F21CA) - Project Lead

January 2025 – April 2025

- Developed a RAG-based medical QA system over 411 NHS records.
- Designed a tri-modal interaction pipeline integrating text, speech, and vision.
- Evaluated retrieval-augmented vs baseline performance using semantic similarity metrics.
- Demonstrated measurable reliability improvement from retrieval grounding (+5.6%).

EXPERIENCE

AI “Learning Buddy” Project Intern – Heriot-Watt University

July 2025 – September 2025

- Built a retrieval-based learning assistant over course materials using LlamaIndex and Milvus.
- Designed ingestion and chunking strategies for retrieval across heterogeneous educational content.
- Analysed embedding and retrieval behaviour to evaluate response quality and relevance.
- Highlights retrieval behaviour across heterogeneous educational domains.

Data Science & Machine Learning Intern – Love For Data

August 2023

- Designed a popularity index across 1,970 Play Store apps spanning US and Pakistan datasets.
- Modelled feature–outcome relationships using a Random Forest regressor to analyse performance drivers.
- Highlights differences in app performance drivers across markets through comparative modelling.

Software Engineering Intern – Folio3

June 2023 – July 2023

- Designed and managed a relational database using MariaDB for a Visitor Management System.
- Integrated database with backend via PHP, enabling data flow across system components.

EDUCATION

Heriot-Watt University – BSc Computer Science (Artificial Intelligence)

First Class Honours, 3x Deputy Principal’s Award (2022/23 – 2024/25)

BlueDot Impact – Technical AI Safety Course (2026) – [Credentials](#)

SKILLS

Languages: Python, Java, C#, C

Tools & Infrastructure: PyTorch, Milvus, Docker, Linux, CUDA, Git, Weights & Biases, RunPod