Mariam Hassan

+44 7468 881107
mariam.ihab.mo@gmail.com
mariam-ihab-mohammed

Summary

AI researcher and **DeepMind Scholar** with experience in foundation models, LLMs, and multimodal learning. Completed an MSc in AI for Biomedicine and Healthcare at UCL. Fluent in **French** (French Baccalauréat) and English. Experience prototyping and adapting model architectures for generative and multimodal systems on large, complex biomedical data (e.g., spatial transcriptomics), optimizing LLM components (e.g., tokenization), and studying mechanistic interpretability.

Education

Sep 2024–Oct 2025 **MSc in Artificial Intelligence for Biomedicine and Healthcare (Distinction)**, *University College London (UCL)*, London, UK

- **DeepMind Scholar:** Full scholarship awarded for academic excellence and research potential.
- Relevant coursework: Probabilistic Modeling, Reinforcement Learning, Statistical NLP, Deep Learning, AI for Biomedicine, Domain-Specific AI.

Sep 2019–Sep 2023 **BSc in Computers and Artificial Intelligence (Honors)**, *Cairo University*, Cairo, Egypt

- Key coursework: Theoretical Foundations of ML, Supervised/Unsupervised Learning, Processing of Formal & Natural Languages, Selected Topics in AI, Generative Adversarial Networks, Brain-Computer Interfacing, Autonomous Multiagent Systems, Algorithms Analysis & Design.

Research Experience

Jan 2025–Oct 2025 **Graduate Research Thesis**, *UCL / The Alan Turing Institute*, London, UK

- Identified a fundamental biological plausibility gap in state-of-the-art models for generating spatial transcriptomics from histology, demonstrating systematic constraint violations and poor global molecular consistency on the HEST-1K breast cancer cohort.
- Developed a biologically-constrained Learning Using Privileged Information (LUPI) framework that integrates matched bulk RNA-seq during training to enforce global molecular consistency while preserving H&E-only inference.
- Designed a multi-objective composite loss combining knowledge distillation, bulk-consistency regularization, spatial coherence, and pathway-level constraints to reduce biologically implausible predictions.
- Built a comprehensive validation pipeline incorporating SSIM, GSVA-based pathway correlation, Moran's I spatial autocorrelation, cell-type deconvolution, and gene-category analysis to assess both predictive performance and biological validity.

Jan 2025–Apr 2025 **NLP Research – Low-Resource Languages**, *UCL*, London, UK

- Investigated static vs. dynamic tokenization strategies for LLaMA-3.1-8B and Gemma-7B in morphologically complex languages.
- Implemented embedding-aware vocabulary extensions reducing Arabic OOV rates by 17.4%.
- Developed automated benchmarking infrastructure evaluating fertility, compression ratios, inference efficiency, and downstream QA/classification performance with LoRA fine-tuning.

Sep 2022–Sep 2023 **Undergraduate Research Thesis**, *Cairo University*, Cairo, Egypt

- Developed a metaheuristic system for the Dynamic Vehicle Routing Problem (DVRP) with real-time map changes and demand uncertainty.
- Built a hybrid model using Fuzzy Receding Horizon Control and Genetic Algorithms for dynamic assignment and route optimization.
- Achieved substantial reductions in simulated distance and waiting times; built a full-stack demonstrator app (Flutter, Dart, Python, Firebase).

Recent Projects

Feb 2026 **AI Meeting-to-Execution Autopilot**, *Cactus x Google DeepMind Hackathon*, London (Global)

- Built an agentic prototype for meeting transcription, task extraction, and execution tracking.
- Achieved #1 London and Top 5 Global ranking (99.5% score; 16ms latency; 100% on-device ratio).

Jan 2026 - Feb 2026 **Mechanistic Interpretability Research Agent**, *Independent Research*

- Building a long-running agent harness for automated interpretability experiments and reproducible pipelines.
- Implements orchestration (function calling, tools, planning, retrieval, sandboxing) for execution and tracking.
- Automates activation/attribution patching with evaluation loops to test robustness and reduce spurious findings.
- Inspired by a BlueDot Impact Project (Mar–Jun 2024): GPT-Neo bracket-closing circuit study.

Professional Experience

May 2021–Aug 2024 **Cloud AI & Data Analytics Consultant**, *Bizbrain*, Cairo, Egypt

- Led end-to-end client engagements through an independent consultancy across multiple long-running contracts.
- Designed and deployed cloud-based AI solutions (Azure AI, AWS ML), optimizing infrastructure and reducing compute costs by 42%.
- Built scalable data pipelines and MLOps workflows (Azure ML, CI/CD, Docker), improving processing speed by 2.5×.
- Developed analytics dashboards and automated reporting systems (Power BI, Tableau), reducing manual workload by 250+ hours/month.

Technical Skills

Programming	Python, C++, C, Java, R, Julia, MATLAB, SQL; Bash
ML/AI	Transformers, Diffusion Models, Multimodal Learning, Few-shot Learning, RL, GANs, Bayesian Methods
Agentic Systems	Function calling, tool use, planning, retrieval, sandboxing, long-running orchestration
NLP	LLMs, attention mechanisms, dynamic tokenization, subword segmentation, cross-lingual transfer, sentiment analysis
Vision	CNNs, object detection (YOLO), image segmentation (U-Net), feature extraction, transfer learning; WSI/medical imaging
Frameworks	PyTorch, TensorFlow, JAX; Hugging Face (Transformers, Tokenizers, Datasets)
Libraries	scikit-learn, NumPy, Pandas, SciPy, Matplotlib, Seaborn, Plotly
Cloud & MLOps	Azure ML, AWS SageMaker; Docker, Kubernetes; MLflow, W&B; Git, CI/CD, DVC
Data	EDA, statistical testing; Tableau, Power BI

Certifications

- Microsoft Certified: Azure AI Engineer Associate
- Azure Responsible AI Workshop Coach
- AWS Certified Cloud Practitioner
- AWS Certified AI Practitioner – Subject Matter Expert Contributor
- AI Safety Fundamentals – BlueDot Impact
- AI Safety, Ethics & Society – Center for AI Safety (CAIS)
- Machine Learning & Deep Learning Specialization – DeepLearning.AI
- AI for Medicine Specialization – DeepLearning.AI