

ADAM EL-KADRI

P: +44 7301 961177 **E:** adamelkadri2@gmail.com
L: linkedin.com/in/adamelkadri **G:** github.com/adamelkadri

Education

Royal Holloway, University of London

BSc Computer Science & Artificial Intelligence

Expected Graduation: June 2027

London, United Kingdom

- **First-Class Honours (Predicted)**
- **Relevant Modules:** Deep Learning, Natural Language Processing, Machine Learning, Intelligent Agents and Multi-Agent Systems, Quantum Computation, Computational Finance, Full Unit Project in Artificial Intelligence

Work Experience

Zenithr

AI Engineer Intern

June 2026 – Present

Hybrid

- Reduced crawler runtime from **138s to 43s** by parallelising a 30-page workload with **Celery**, **Redis**, and **4 concurrent workers**.
- Built an **LLM-based extraction** pipeline using **Anthropic Claude** and **Playwright** to parse candidate profiles from web pages and **PDFs**, achieving a **94.7%** crawl-job success rate.
- Developed **31 REST endpoints** across an **8-service Docker** backend with **FastAPI**, **PostgreSQL**, **Redis**, **AWS S3**, **158 tests**, **CI/CD**, and **OAuth/OIDC**.

Royal Holloway, University of London

Research Assistant

June 2026 – Present

London, United Kingdom

- Fine-tuned **Qwen2.5-7B-Instruct** using **QLoRA** on the **SocSci210** behavioural dataset, training on **200 instruction-response examples** over **1 epoch** and reducing training loss from approximately **1.84** to **0.98**.
- Engineered a reproducible **Python** experimentation workflow using **Pandas**, **Hugging Face**, and **LLaMA-Factory** to clean, sample, and format behavioural data across **210 social science studies**.
- Benchmarked base and **QLoRA-adapted** model outputs using held-out behavioural examples, evaluating generalisation across **170 seen studies** and **40 held-out studies** through exact-match accuracy and reconstruction fidelity.

Projects

Time Series Forecast AI Copilot | *Website*

LangChain | **LangGraph** | **Streamlit** | **SQLite**

- Built a full-stack **AI analytics copilot** that translates natural-language queries into executable **SQL** and **time-series forecasts**, enabling users to analyse uploaded **CSV/Excel** datasets without writing code.
- Engineered a multi-agent **LangGraph** workflow powered by the **Anthropic Claude API**, integrating **LLM reasoning**, query generation, forecasting logic, and per-session **SQLite** isolation.
- Implemented an asynchronous data pipeline with structured error handling for invalid queries, model failures, and cloud deployment using thread-safe async execution.

ECG Arrhythmia Classifier | *Source Code*

PyTorch | **FastAPI** | **Docker**

- Built an end-to-end **1D CNN** with approximately **44K parameters** in **PyTorch** to classify **5 arrhythmia types** from **109K+ ECG heartbeats**, achieving **93.8% accuracy** and **0.77 macro-F1** on **21,892 held-out beats**.
- Handled **113× class imbalance** using **weighted loss**, achieving **74–93% recall** on rare classes, then deployed the model via a **Dockerised FastAPI REST API** with a **17-test CI pipeline**.

Technical Skills

Languages: Python, Java, JavaScript, TypeScript, C, SQL

AI/ML: PyTorch, TensorFlow, Scikit-learn, Pandas, NumPy, Hugging Face, QLoRA, LLaMA-Factory, CNNs, NLP, LLMs, Time Series Forecasting, Model Evaluation, Data Preprocessing, Classification

LLM & Agent Frameworks: LangChain, LangGraph, Anthropic Claude API, LLM Extraction, RAG, Prompt Engineering

Backend & Full-Stack: FastAPI, React, Node.js, Express.js, Spring Boot, REST APIs, Streamlit, OAuth/OIDC

Tools & Databases: Docker, Git, GitHub Actions, CI/CD, Linux, PostgreSQL, SQLite, Redis, Celery, Playwright, AWS S3, Supabase, pgvector