Moiz Imran

Greater London, UK linkedin.com/in/moizimran github.com/Moiz-I

EDUCATION

MEng Computer Science 09/2022 – 05/2026
University College London London, UK, Grade: 1st (expected)

- Experience: Teaching Assistant (Object-Oriented Programming Module); assisted in lab sessions, helping students debug and answering any questions
- Notable Project: Worked with Microsoft, IBM, NHS and the UCL CS department on a touchless computing project
- Notable Modules: Algorithms, Discrete Maths, Logic, Theory of Computation, Software Engineering, OOP, Functional Programming, Computer Systems

AI Startup School 07/2025 – 07/2025
Y Combinator San Francisco

- Along with 2,000 of the world's smartest students, learned from the greatest minds in tech and AI, from founders to researchers.
- Everything from discussing gene editing in a café to meeting with top YC founders, this experience was as valuable as my degree, if not more.

PROFESSIONAL EXPERIENCE

Forward Deployed Engineer 11/2026 – Present
Arrakis London, UK
Applied AI in the physical economy

Software Engineer Intern 07/2025 – 09/2025
HubSpot London, UK

- Worked on a core fintech team building internal quoting infrastructure powering 1000+ daily enterprise deals. Focus on accuracy and real-time sync across billing + CRM systems
- Collaborated cross-functionally with legal, sales, and billing to model complex discounting and contract logic in a maintainable & testable way
- Identified and fixed a concurrency bug in the approval routing system that caused intermittent delays for large-volume quotes (>500 line items)
- Working on a performance optimisation that could reduce quote generation latency by 40% through optimised access patterns and strategic caching
-

Software Engineer Intern 07/2023 – 10/2023
IMMIX London, UK

- Shipped production-ready tools supporting digital asset trading across multiple exchanges, ensuring seamless failover and low-latency reliability (<30µs)
- Built internal dashboards for monitoring market connectivity and order health using NextJS, and websockets, reducing time-to-detection for live incidents
- Helped deliver a fault-tolerant, 24/7 trading system with privacy-focused, self-hosted solution
- Delivered key front-facing features, collaborating across teams and reporting to senior management
- Took initiative by building and launching an AI documentation chatbot for clients

SKILLS

Languages & Core Technologies
Python, Java, C++, TypeScript, Docker, AWS

Specialised Skills
Abstract Syntax Tree Manipulation, Advanced Image Processing, ML (Tensor, OpenCV), Data Analysis (NumPy, Pandas), Technical Documentation

PROJECTS

UCL MotionInput 10/2023 – 05/2024

- Worked with Microsoft, IBM, NHS and the UCL CS department on a touchless computing project

- Utilised open-source ML models and engineered a robust solution for potential deployment in Great Ormond Street Hospital.
-

Palmy 2024 – Present
A side-project with friends that's reached over 300k users since launch. We bring AI to *actually* help students (not just a wrapper). Features include topic mastery, infinite test mode (dynamically adjusting questions) and recently a Cursor-like agent for your notes.

Council of Claude 06/2026
Open-Source Multi-Agent Deliberation System
Built a system that eliminates single-model bias by running parallel LLM queries with assigned personas, then forcing anonymous cross-pollination before synthesis, solving the problem of AI "groupthink" in high-stakes decisions.
- Designed a convergence protocol that detects genuine agreement vs politeness bias between models
- Single-file async architecture for zero-friction installation across AI coding agents
- Published open-source; adoptable in one command

Personal AI systems lab 01/2026 – Present
(Work in Progress)
- Self-hosted a multi-profile AI agent stack on EC2 using Hermes (Nous Research), with isolated Telegram/Discord agents for conversation, capture, and engineering workflows, plus cron routines, skill loading, session isolation, and layered context management.
- Built an AI resource pipeline using Telegram ingestion, FastAPI, PostgreSQL/RDS, Gemini embeddings, LLM enrichment, semantic clustering, and a Vercel triage UI to turn saved resources into structured project knowledge.
- Built a blind human-vs-model evaluation arena with parallel model calls, response standardisation, anonymised LLM judging, rubric scoring, performance history, and adaptive difficulty.

The Correct Answer Trap, MEng thesis, UCL · supervised by Sahan Bulathwela 2025 – 2026
Published in AIED 2026 (Springer LNCS) + PEAF workshop
- Found that AI tutors miss misconceptions when a student's flawed reasoning still reaches the right answer: across 21k real student responses, fine-tuned BERT hit 98.9% overall accuracy but caught only 57% of these cases.
- Tested prompted LLMs (the best caught 84%, with ~4 false alarms per real catch) and designed a pipeline that detects, verifies and escalates uncertain cases.

AWARDS

1st Place – Optiver Challenge, UCL x LSE Data Science Hackathon
Built the top-ranked match-making algorithm on Optiver's mock trading platform, achieving highest PnL through historical data analysis and strategic order flow logic.

Buildspace 'nights & weekends' Hackathon
Their flagship event, had 100,000 participants working on bringing their ideas to life.
I worked on a game from scratch, posted my updates on social media, getting 2000 people on my waitlist, pitching to strangers, and even making it open-source to inspire others to make their own game.
I was chosen to be one of the 20 final winners, earning a $5K grant to continue working on my idea.

Arc's 'Take Back the Internet' Hackathon
The Browser Company New York
I used APIs from Arc and OpenAI and design to create a seamless extension for their browser, allowing search results in Google to also contain a LLM-summarised answer to the user's query.
Ranked first in the judges' categories of creativity, utlity, technical achievement.
(This was launched 2 years before Google launched their "AI Summaries")