Abdul Moiz Ahmed

Lahore, Pakistan | 0328-4638346 | l232528@lhr.nu.edu.pk
linkedin.com/in/abdul-moiz-ahmed-04b38824a | github.com/Moiz2003

OBJECTIVE

AI & Software Engineer focused on applied machine learning, multi-agent LLM systems, and production backend architecture. Experienced in building, deploying, and stress-testing real systems—from XGBoost forecasting pipelines to adversarial multi-agent safety loops—with an emphasis on reliability, fallback design, and honest engineering.

EDUCATION

FAST NUCES Lahore, Pakistan
Bachelor of Computer Science (6th Semester) *2023 – Expected 2027*
Relevant Coursework: Artificial Intelligence, Computer Architecture, Software Testing & QA Automation, Design & Analysis of Algorithms

TECHNICAL SKILLS

Machine Learning & AI: XGBoost, scikit-learn, time-series feature engineering, multi-agent LLM orchestration, RAG concepts, prompt design, LLM APIs (OpenAI-compatible, DeepSeek, Gemini, Ollama / local vLLM)
Languages: Python, JavaScript/TypeScript, Java, C++, Swift, Kotlin
Backend & Systems: FastAPI, Node.js/Express, async concurrency (asyncio), WebSockets, REST APIs, MongoDB (Motor/Mongoose), InfluxDB, JWT/OAuth
Cloud & Infrastructure: AWS EC2, Docker / docker-compose, Nginx, PM2, Vercel, CI/CD (GitHub Actions)
Frontend & QA: React, Vite, Tailwind, Recharts/Nivo; Playwright, Cypress, Karate DSL, Gatling, JUnit 5

EXPERIENCE

Lead Engineer Jan 2024 – Present
Gynae Care Center (Clinify Project) – www.gynaecarecenter.com *Lahore, Pakistan*

- Built and deployed a multi-tenant healthcare SaaS on AWS EC2 (Nginx/PM2) with a MongoDB Atlas backend, serving a live gynaecology clinic.
- Engineered secure REST APIs with role-based JWT auth and Google OAuth 2.0, enforcing rate limiting, Helmet.js headers, and NoSQL/XSS sanitization.
- Integrated Agora SDK for token-authenticated WebRTC video consultations and automated medical notifications via Nodemailer.
- Optimized the React/Vite frontend to a sub-0.5MB bundle using React.lazy and Suspense for heavy charting components.

AI & ENGINEERING PROJECTS

Hyperion — Adversarial Multi-Agent Diagnostic Pipeline | ***Node.js, Express 5, LLM Vision APIs, MongoDB*** *LabLab Hack*

- Built a multi-agent Vision → Drafter → Critic pipeline where an adversarial "Critic" agent bias-rejects unsafe or incomplete clinical drafts, driving a bounded consensus revision loop.
- Implemented SHA-256 image hashing with MongoDB result caching (24h TTL) and SSE streaming with per-agent event emission and partial-state recovery on timeout.

Materna — Safety-First Obstetric Decision-Support System | ***Python, FastAPI, Pydantic, Multi-Agent Orchestration***

- Architected an orchestrator coordinating four specialist agents (intake, dating-risk, guideline, auditor) over typed Pydantic contracts, with a deterministic non-LLM safety gate enforcing human-in-the-loop escalation.
- Engineered a hash-chained, tamper-evident audit trail with a verification endpoint, plus graceful degradation when LLM/vision providers are unavailable.
- Wrote an adversarial hardening test suite covering gate-bypass, prompt-injection resistance, determinism, and audit-tamper detection.

GridLock OS — Predictive Energy Control Engine | ***Python, XGBoost, InfluxDB, Docker***

- Trained an XGBoost battery/SoC forecasting model using time-aware feature engineering and chronological train/test splits, with live inference from a rolling InfluxDB telemetry window.
- Built a resilient 5-minute autonomous control loop integrating inverter (Modbus) and HVAC telemetry to throttle loads and extend daily battery runtime, with retries, stale-data tagging, and degraded-mode fallbacks.

KAIROS — AI-Assisted Crypto Trading Engine | ***Python, FastAPI, asyncio, WebSockets, MongoDB***

- Built a human-in-the-loop trading engine combining technical-analysis signals (RSI/EMA) with an LLM sentiment "reality check" and a risk gate before execution.
- Engineered a single-upstream Binance WebSocket fan-out with exponential-backoff reconnection, async I/O offloading (asyncio.to_thread), SHA-256 request dedup caching, and Decimal-precise order sizing against exchange LOT_SIZE filters.