

Mohammed Zuhair Hussain

AI Engineer @ Teclusion AI | Prev. AI Intern @ PwC | GenAI, LLMs & ML Systems
mohammedzuhairhussain28@gmail.com | +91 9701007757 | [GitHub](#) | [LinkedIn](#) | [Portfolio](#) | Hyderabad, India

Education

B.Tech in Artificial Intelligence Engineering, Mahindra University, Hyderabad Sept 2022 – Aug 2026
CGPA: 8.13/10.0 | Final-semester Industry Internship/Project graded **A+ (10/10)** | Final-year SGPA: 8.79, 10.0
Finalist — Synchrony UniHack 2025 | Academic **merit scholarship** worth INR 100,000 for excellent academic performance

Experience

AI/ML Engineer Intern — Teclusion AI

September 2025 – Present

- Co-built **Visentix**, a multi-tenant privacy-intelligence SaaS, as **one of two engineers**: an end-to-end **LLM pipeline** (SSRF-protected intake, clause decomposition, LLM classification, dynamic benchmarking, deterministic scoring, frozen report snapshots) on **FastAPI** and **PostgreSQL (Supabase + pgvector)** with a self-hosted **Qwen3-8B served via Ollama**; deployed live across **Azure** (Dockerized API behind Caddy), a **RunPod GPU inference pod** over a private Tailscale mesh, and Cloudflare Workers.
- Enforced strict **LLM/arithmetic separation**: the LLM only classifies clauses into **eight privacy domains** (~1.6 s/clause, **384-dim sentence-transformer embeddings** for similarity search), while **14 versioned deterministic formulas** compute every score with full **source-to-score lineage**, **confidence-based suppression** of low-confidence outputs, and **byte-identical reproducible PDF reports** (SHA-256 regression-tested).
- Built **human-in-the-loop and safety layers**: an SME review workbench (confirm/edit/dismiss AI findings, **PII detection and redaction**) feeding a **training-label flywheel**, plus a deterministic **guardrail filter** that hard-fails any output containing banned legal-verdict language.
- Hardened the platform for **multi-tenant isolation**: PostgreSQL **row-level security** with JWT-scoped access so no tenant can read another's documents or reports, plus rate limiting and graceful fallbacks on the inference path.
- Shipped **BluePearl**, an Intelligent Document Processing platform that lets financial teams ask one question across hundreds of reports and get accurate, source-cited answers — running entirely on locally hosted LLMs so client data never leaves their infrastructure, and cutting manual document review time by **90%**.
- Built **Vantage**, a Dynamic Pricing Engine combining **ARIMA/SARIMA** time-series forecasting with ML models to predict demand and price elasticity for every SKU, reacting to competitor moves within business-set margin and brand-price floors — driving a **6–10%** increase in optimized revenue and a **50% reduction** in manual pricing effort.

AI Intern, Emerging Tech — PricewaterhouseCoopers (PwC), Hyderabad, India

Jun 2025 – Aug 2025

- Developed a full-stack **AI financial-advisory platform** (FastAPI, Next.js, Firebase): an agentic system on Google Gemini grounding personalised investment and tax guidance in a custom knowledge graph, answering in **under 2 seconds**.
- Co-authored an internal **white paper** (draft) on AI-enabled software development, driving research and framework design across LangChain, agentic AI, and generative AI in the SDLC.
- Prototyped and demoed AI solutions for stakeholders, carrying features from concept to shipped implementation in an enterprise consulting environment.

Projects

Aperture — Cash-Flow Credit Decisioning Platform

[GitHub](#)

- Built end-to-end a credit-decisioning platform for **applicants banks cannot score**: underwrites from verified bank-transaction evidence (Account Aggregator) rather than a bureau score, served via FastAPI REST APIs with a React + TypeScript reviewer dashboard — a **9-gate deterministic policy engine** makes every call, with **zero AI in the decision path**.
- Developed a transparent **11-feature cash-flow scorecard** (scikit-learn) exposing each feature's exact contribution to every decision, benchmarked against calibrated XGBoost — **ROC-AUC 0.787, KS 0.45** vs a 0.726 logistic baseline.
- Made decisions auditable and fair by design: point-in-time features with full lineage in PostgreSQL, a **fairness allow-list** blocking protected attributes at build time, 8 deterministic fraud detectors, and a **hash-chained audit ledger**.
- Added two grounded AI assistants (“explain this decision”, “ask the codebase”) that phrase only verified facts and fall back to templates rather than guess — backed by **393 automated tests** and strict static typing.
- Packaged the stack for reproducible deployment — Dockerized services on AWS EC2, with pytest and **mypy (strict)** gating every change so any past decision can be replayed from the same commit.

Technologies: Python, FastAPI, PostgreSQL, scikit-learn, XGBoost, React, TypeScript, Docker on AWS EC2

Time Series Forecasting of Stock Prices — Model Benchmarking Study

[GitHub](#)

- Benchmarked **five model families** (ARIMA, GARCH, Prophet, KNN, feed-forward networks in TensorFlow) for **30-day stock-price forecasting**, weighing accuracy (RMSE, MAPE) against complexity in a reproducible experiment pipeline.
- Grounded every comparison in statistical validation — ADF stationarity tests, ACF/PACF analysis, Ljung-Box diagnostics — with **GARCH volatility modelling** for risk-aware forecasts.

Technologies: Python, TensorFlow, Statsmodels, Prophet, scikit-learn, Pandas

WISP — Offline Mesh Messaging PWA

[GitHub](#)

- Architected a **serverless Progressive Web App** for crisis communication across three transport tiers (WebRTC over Wi-Fi, animated QR sync, printed posters), deployable via URL in **under 30 seconds** with zero backend — offline-first, with IndexedDB persistence, TTL-based auto-cleanup, and a live force-directed mesh visualization.
- Engineered a gossip-style sync engine with Bloom-filter anti-entropy, Ed25519 signatures, and **NaCl end-to-end encryption**, achieving **sub-200 ms LAN sync** with chain-of-custody hop stamping and panic-wipe identity rotation in a 180 KB gzipped bundle.

Technologies: React 18, TypeScript, WebRTC, IndexedDB (Dexie), Ed25519, TweetNaCl, Vite, Service Workers

Skills

Core Areas: Machine Learning, Deep Learning, NLP, Computer Vision, Generative AI, LLM Applications
Languages: Python, SQL, TypeScript/JavaScript, Java, C
ML & Deep Learning: PyTorch, TensorFlow, Keras, scikit-learn, XGBoost, time-series forecasting (ARIMA, SARIMA, GARCH, Prophet), Statsmodels, NumPy, Pandas, OpenCV, model monitoring (calibration, drift)
LLM & GenAI: RAG, embeddings & vector search (pgvector, sentence-transformers), LangChain, LangGraph, Hugging Face, Ollama (local LLM serving), agentic systems, prompt engineering, LLM guardrails & evaluation
Backend & Web: FastAPI, REST APIs, async I/O, Node.js, Express.js, React, Next.js, Streamlit
Data & Cloud: PostgreSQL (Supabase, pgvector), MongoDB, Firebase, Docker, AWS (EC2), Azure, Vercel, Linux, Git/GitHub
Engineering Practices: multi-tenant security (RLS, JWT), rate limiting & fallbacks, pytest, mypy (strict), PII detection & redaction

Key Courses Taken

- **Machine Learning & AI:** Machine Learning and Python[#], Deep Neural Networks, Applied AI for Industry[#], Computational Sequence Modelling[#], Visual Data Analytics[#], Optimization Techniques for AI[#], Foundations of NLP, Artificial & Computational Intelligence, Digital Image Processing, Computational Genomics
- **Computer Science:** Data Structures, Design and Analysis of Algorithms, Object Oriented Programming[#], Operating Systems, Database Management Systems, Software Engineering, Computer Networks, Theory of Computing
- **Mathematics:** Mathematics I–III, Numerical Methods, Discrete Mathematical Structures, Signals & Systems

[#] Grade A or above

Positions of Responsibility

General Secretary — Mahindra University Jun 2024 – Jun 2025

- Acted as the key liaison between **3,000+ students** and university management, resolving operational issues and driving collaborative problem-solving during events.
- **Led a team of 16** to organise 5 major campus events — Farewell, Orientation, Freshers, Convocation (for **1,000+ graduates**), and the cultural event “RRR” — owning planning, delegation, and on-ground execution.
- Expanded club activities and systemized the administrative structure, and contributed to the initial planning of **AEON**, the university’s first-ever Tech Fest.