



# Oluwamayowa T. Adeoni

✉ [oluwamayowaadeoni@gmail.com](mailto:oluwamayowaadeoni@gmail.com)  [github.com/pope001](https://github.com/pope001)  [linkedin.com/in/mayowa-adeoni](https://www.linkedin.com/in/mayowa-adeoni)  
<https://mayowaadeoni.com/>

---

## PROFESSIONAL SUMMARY

Fullstack AI/ML Engineer specializing in Agentic AI and Multi-Agent Orchestration. Expert in building production-grade systems using LangGraph, MCP, and RAG architectures. Proven track record in deploying autonomous "Human-in-the-Loop" systems that solve complex search and transactional challenges with 99% schema compliance.

---

## EDUCATION

**Obafemi Awolowo University**, Nigeria — *B.Sc. Computer Science with Mathematics*

*August 2017 - July 2023*

- Relevant Coursework: Data Structures and Analysis of Algorithms - **CSC 315**, Introduction to Artificial Intelligence - **CPE 316**, Automata Theory & Computability - **CSC 317**, Techniques in Software Development - **CSC 407**, Natural Learning Processing & Application - **CPE510**
- 

## EXPERIENCE/PROJECTS

**HeySynth | Remote (Austin, USA) — Fullstack AI/ML Engineer**

*July 2025 – Present*

*Architecting the core multi-agent operating system empowering CPG brands to autonomous operations.*

- Architected a core multi-agent operating system for CPG brands, utilizing LangGraph to build 4 autonomous agents that reduced operational overhead by 70%.
- Built and enhanced an end-to-end Python ML forecasting service for retail demand, implementing tiered model selection (NeuralProphet, SES, Holt) and fallback strategies for sparse datasets.
- Developed event-driven architectures using FastAPI and Google Pub/Sub, integrating asynchronous training/prediction workers and ensuring sub-200ms latency for agent responses.
- Improved production reliability by implementing conservative/aggressive forecast modes and engineering cyclic graph state-machines for "Human-in-the-Loop" validation with 99% schema compliance.
- Designed multi-agent reasoning workflows using Chain of Thought (CoT) and MCP architectural logic to ensure standardized data injection and eliminate hallucinations in inventory workflows.
- Led the fullstack deployment stack across Google Cloud (Vertex AI), MongoDB, and React/TypeScript, visualizing real-time agent "thought chains" for complex transactional workflows.

**InspireEdge | Remote (UK) — Fullstack AI/ML Engineer**

*April 2025 – July 2025*

- Architected and deployed production AI modules for e-commerce SMEs, including market intelligence and risk detection systems, enabling merchants to transform raw data into actionable growth signals.
- Engineered end-to-end LLM-powered workflows (summaries, insights, and recommendations) using precision prompt design and model selection to deliver explainable, high-fidelity outputs for non-technical users.
- Engineered end-to-end LLM-powered workflows for market intelligence using precision prompt design and RAG-based architectures to deliver explainable, high-fidelity insights
- Managed fullstack integrations between backend AI services (FastAPI) and major e-commerce platforms like Shopify and WooCommerce, ensuring seamless delivery of behavioral insights to merchant dashboards.
- Optimized e-commerce modules for cart-abandonment prediction and anomaly detection, focusing on building practical, revenue-impacting intelligence that directly reduced operational risk for users.
- Benchmarked and monitored model performance using Weights & Biases (W&B), ensuring that production-grade forecasting and behavioral models remained accurate and scalable in low-latency environments.

**Nabafat.AI Technologies | Remote (California, USA) — AI/ML Engineer**

*January 2024 - March 2025*

- Architected and deployed multi-agent reasoning workflows using LangChain for enterprise-scale applications, including a "Startup Adviser" and "Personalized Learning Agent" that automate complex decision-making through structured agentic planning.
- Engineered a multi-LLM orchestration pipeline integrating OpenAI GPT-4, Llama 3.1 (Groq), and DeepSeek to build a sophisticated AI coding assistant; implemented intelligent task routing that optimized token usage and response accuracy.
- Developed "NeuroScan," a production-ready medical imaging platform utilizing TensorFlow Lite to classify brain tumors from MRI scans, delivering instant diagnostic confidence scores and heatmaps for clinical interpretability.
- Built an end-to-end cybersecurity solution utilizing XGBoost to detect phishing attempts with 99% accuracy, handling the entire lifecycle from data preprocessing to deployment as a low-latency REST API.
- Developed fullstack AI architectures by integrating robust Python (FastAPI/Flask) backend services with interactive React and Streamlit interfaces, providing real-time visualization of agent "thought chains" and reasoning logs.
- Containerized and scaled AI applications using Docker, ensuring reproducible deployment workflows and high availability across cloud-native environments for diverse SaaS products.

### **Quantum Leap Limited | Remote— AI/ML Engineer**

*March 2023 - August 2023*

- Optimized training pipelines and inference economics by transitioning core TensorFlow models to JAX and NumPy, resulting in a 30% reduction in cloud infrastructure costs.
- Developed and trained high-accuracy models for sentiment analysis and customer churn prediction using TensorFlow to drive data-driven business decisions.
- Streamlined MLOps workflows by implementing MLFlow for robust experiment tracking and data versioning, ensuring reproducible results across training cycles.

---

## SKILLS

**AI/ML & Agentic:** LangGraph, LangChain, Retrieval-Augmented Generation (RAG), Chain-of-Thought (CoT), Model Context Protocol (MCP), TensorFlow, PyTorch, JAX/Flax, Vertex AI, Gemini Pro, Llama 3, XGBoost, Scikit-Learn, CoT reasoning, OpenAI Agent SDK.

**Fullstack & Data:** Python, TypeScript/React, Next.js, Node.js, FastAPI, Flask, Django, SQL, NoSQL (MongoDB), Vector DBs (ChromaDB).

**Infrastructure & MLOps:** AWS, Google Cloud (GCP), Docker, Kubernetes, Weights & Biases (W&B), MLFlow, Git, Linux.

**Testing & Ops:** Weights & Biases (W&B), MLFlow, Docker, CI/CD, A/B Testing logic.

---

## LEADERSHIP/MEMBERSHIP

- Aug 2017 - July 2023 | Member - Google Student Developer Club.
- Jun 2020- July 2023 | Member, BlockchainOAU Student Club.
- 2017- Present | Member, DataScienceNigeria(DSN).