

# Moses Adebayo

+234 812 067 3365 | [adebayomoses597@gmail.com](mailto:adebayomoses597@gmail.com) | [linkedin.com/in/adebayo-moses](https://linkedin.com/in/adebayo-moses) | [mosesadebayo.com](https://mosesadebayo.com) | [github.com/mosessmax](https://github.com/mosessmax)

## EDUCATION

---

### Lagos State Polytechnic

Higher National Diploma in Computer Science

Lagos, NGA

May 2022 – Sept 2024

## EXPERIENCE

---

### AI Backend Engineer

June 2025 – Present

MIVA Open University

Abuja, NGA

- Architected and lead development of MIND 2.0, an AI intelligence layer embedded in the Miva Campus LMS, designing a modular RAG pipeline using LangChain, LangGraph, and Qdrant to power context-aware tutoring, adaptive assessments, and personalized learning pathways.
- Built and deployed LLM orchestration infrastructure with FastAPI, implementing multi-step agentic workflows, prompt versioning, and structured output validation to serve AI-driven features reliably across 25,000+ active learners.
- Designed a real-time voice AI interview sparring tool (Sparr) using LiveKit Agents SDK, integrating speech-to-text, RAG-based context retrieval, and LLM response generation within a low-latency streaming pipeline.
- Instrumented end-to-end LLM observability with Langfuse, tracking trace quality, retrieval relevance, latency, and token costs across AI services to enable continuous evaluation and optimization of model performance.
- Defined data contracts and integration architecture between AI services and the LMS engineering team, negotiating API boundaries for the phased AI roadmap and maintaining high-traffic endpoints.

### Full-Stack Developer

Nov 2024 – May 2025

JOPHAB

Remote

- Architected and built production-grade backend services using Node.js, TypeScript, and Express.js, implementing RESTful APIs with Drizzle ORM for type-safe database operations and achieving 40% faster response times compared to the previous architecture.
- Designed scalable PostgreSQL database architecture with advanced normalization, complex relationships, and optimized indexing patterns, supporting high-traffic applications processing 1M+ daily transactions.
- Led system design for microservices infrastructure on AWS (EC2, S3), implementing event-driven communication patterns and horizontal scaling strategies with PM2 and Nginx load balancing.
- Established CI/CD pipelines with Docker containerization and automated testing workflows, enabling zero-downtime multi-environment deployments and reducing deployment errors significantly.

### Frontend Lead Developer

Oct 2024 – Nov 2024

Foodpulse

Aberdeen, United Kingdom

- Led frontend development while contributing to backend services with Express.js and RESTful APIs, implementing type-safe API consumption patterns with TypeScript and comprehensive error handling across the platform.
- Implemented advanced state management and server-side data fetching with Tanstack React Query, reducing API calls by 45% through optimized caching and real-time updates via WebSocket connections.

### Full Stack Developer

Aug 2022 – Jan 2024

Remotual

Illinois, Chicago

- Designed and built enterprise-scale backend infrastructure with Node.js, Express, TypeScript, and PostgreSQL, architecting microservices with clear service boundaries and event-driven patterns across distributed systems.
- Architected and deployed containerized microservices on AWS (ECS, RDS, ElastiCache), implementing horizontal scaling and fault-tolerant patterns to handle 10x traffic growth on different client services.
- Built authentication and authorization systems with JWT, OAuth2, and RBAC, implementing granular permission hierarchies and audit logging across multiple client services.

## PROJECTS

---

- MIND 2.0** | *Python, FastAPI, LangChain, LangGraph, Qdrant, PostgreSQL, Redis, Langfuse* 2025
- Designed the full AI architecture for Miva's campus LMS intelligence layer: RAG pipelines with semantic chunking and hybrid retrieval, multi-step LangGraph workflows for adaptive learning, and an agentic tutoring engine with structured output and fallback handling.
  - Built LLM evaluation infrastructure using Langfuse to monitor retrieval relevance, hallucination rates, and latency across production AI services, feeding metrics back into prompt and retrieval optimization cycles.
- Sparr** | *Python, FastAPI, LiveKit Agents SDK, LangChain, Qdrant, Langfuse, React, Vite* 2025
- Built a real-time voice AI interview sparring tool using LiveKit Agents SDK, orchestrating STT, RAG-based context injection, and LLM response generation within a low-latency audio streaming pipeline.
  - Architected as a monorepo with a FastAPI backend handling agent lifecycle, session state, and Qdrant-backed knowledge retrieval, and a React/Vite frontend for live session management.
- Moniva** | *TypeScript, React, Node.js, PostgreSQL* 2025
- Built a visual database management platform ("Figma for databases") enabling developers to design, visualize, and manage relational schemas through an intuitive drag-and-drop interface.
  - Gained significant organic traction on social media, demonstrating clear product-market fit for developer tooling in the African tech ecosystem.
- Papyr** | *Python, FastAPI, PostgreSQL, Redis, React, Vite* 2025
- Architected an AI research paper aggregator pulling from Arxiv, Anthropic, OpenAI, DeepMind, and HuggingFace into a unified, searchable feed with background job processing via Redis-backed task queues.
  - Designed a clean read-later and tagging system on top of a normalized PostgreSQL schema, enabling semantic filtering and prioritization across thousands of ingested papers.
- Docksmith CLI** | *TypeScript, Node.js, Docker, AWS* 2025
- Architected a CLI tool with intelligent framework detection supporting Node.js, Express, Fastify, Next.js, and Python frameworks, generating optimized multi-stage Dockerfiles with cloud-specific configurations for AWS and GCP.
  - Designed an extensible plugin architecture with dependency injection patterns, enabling 15+ framework integrations with 95% detection accuracy.

## TECHNICAL SKILLS

---

**AI & ML:** LLM Orchestration, RAG Systems, LangChain, LangGraph, Agentic Workflows, Prompt Engineering, Vector Search, Embedding Pipelines, LLM Evaluation & Observability

**AI Infrastructure:** Qdrant, Langfuse, LiveKit Agents SDK, Celery, Redis, Background Job Processing, Streaming Pipelines

**Programming Languages:** Python, TypeScript, JavaScript (ES6+), Node.js

**Backend Frameworks:** FastAPI, Flask, Express.js, Fastify, RESTful API Design, Microservices Architecture

**Databases & ORMs:** PostgreSQL, Drizzle ORM, MongoDB, Redis, MySQL, SQL Optimization, Database Design & Normalization

**System Design:** Service Boundaries, Data Flows, Scaling Strategies, Event-Driven Architecture, Caching Patterns

**Cloud & DevOps:** AWS (EC2, ECS, RDS, S3, Lambda, ElastiCache), Kafka, Docker, Kubernetes, CI/CD Pipelines, Nginx, PM2, Cloudflare

**Security:** JWT, OAuth2, RBAC, Session Management, Data Access Controls

**Frontend:** React, Astro, Vite, Tailwind CSS, Framer Motion, TypeScript, State Management, Server-Side Rendering

**Tools & Monitoring:** Git, Postman, Swagger, Logfire, Grafana, Flower, Sanity CMS