Odunukan Abdulhameed

AI Systems Engineer / Backend Engineer (LLM Platforms)

odunukanabdulhameed@gmail.com +2347065393688

odunukan-abdulhameed-main-port.vercel.app/ linkedin.com/in/abdulhameed-odunukan/

github.com/Abdulhameed556 https://medium.com/@odunukanabdulhameed

https://huggingface.co/Hameed13

Professional Summary

Backend & AI Systems Engineer specializing in building scalable backend platforms and LLM-powered systems. Experienced in AI agent orchestration, multi-platform data integration, and secure API architecture. Proven track record of designing intelligent systems that unify heterogeneous data sources through conversational AI, while optimizing cost, performance, and reliability in production environments.

Education

Bachelor of Engineering in Electrical and Electronics Engineering, 2018 – 2024
Federal University of Agriculture, Abeokuta Ogun State, Nigeria

Professional Experience

AI Engineer, *Mactay* 01/2026

- Architected the AISA Omni-Channel Ecosystem by engineering a suite of conversational AI and voice agents leveraging RAG and intelligent pause-handling to automate customer engagement across web platforms and LMS website.
- Architected a production-grade Agentic Email Intelligence System for Gmail and Outlook, featuring a multi-format knowledge base (PDF, DOCX, CSV), semantic RAG retrieval over pgvector, and an LLM pipeline that classifies intent, urgency, sentiment, and routes emails to teams with confidence scoring. Built on a hexagonal architecture with an async job queue supporting both auto-reply and human-in-the-loop modes. Delivered for **First Bank and Altbank.**
- Developed a Multilingual AI QA Pipeline to analyze and score 5,000+ calls for sentiment and responsiveness while integrating cloud storage from OneDrive, Google Drive, and SharePoint and expanding support to French, Portuguese, Chinese, and Spanish.
- Built a Dual-Interface Call Centre Support System featuring an agent-facing chat interface for real-time customer interaction assistance and a dedicated admin control panel for knowledge base management and monitoring.
- Built **AISA Whispar**, a real-time AI voice copilot for Alt Bank's call centre agents, combining Deepgram Nova-3 and Google Chirp 3 for live multilingual transcription (English, Yoruba, Hausa) across Twilio, 3CX, and browser channels, with a RAG-powered suggestion pipeline that surfaces contextual agent cues mid-call using pgvector semantic retrieval and an async worker queue.

Backend Engineer (Python) | AI Engineer, *Dex, Afriex* 11/2025 – Present
United State

- Architected SQL Copilot and a custom LLM orchestration framework, enabling natural language querying with multi-step reasoning across PostgreSQL, MongoDB, Redshift, and Google Sheets.
- Engineered a unified data layer with API connectors for Shopify, Stripe, and Slack via GraphQL and REST, with end-to-end OAuth 2.0 backend workflows — co-led the strategy that landed us #3 Product of the Day on Product Hunt.
- Built an Omni-Channel AI Support Agent for Afriex across chat and voice channels, implementing a four-stage pipeline (input guardrails → scenario routing → RAG + LangGraph orchestration → output guardrails), automated ticket generation with AI tag classification, and DeepEval-powered RAG evaluation for continuous knowledge base monitoring.
- Built a support agent across chat and voice channels implementing a four-stage pipeline (input guardrails → scenario routing → RAG + LangGraph orchestration → output guardrails) for the customer care agents.

Data Scientist, *Rater X* 04/2025 – 11/2025

- Spearheaded the engineering of high-fidelity multilingual datasets for African languages—including Yoruba, Hausa, and Pidgin—scaling production to 100,000+ labeled entries to power enterprise-grade sentiment analysis and voice intelligence systems.
- Architected automated data infrastructure and NLP pipelines that achieved 95%+ accuracy across healthcare and finance domains, utilizing Python and Hugging Face Transformers to deliver production-ready endpoints for the "Sabi AI" product suite.
- Led cross-functional engineering and annotation teams to implement rigorous quality control and monitoring frameworks, ensuring dataset integrity and cultural relevance while optimizing FastAPI-based inference APIs for real-time, distributed workflows.

AI/Machine Learning Engineer, *Nexus Mind Group & COLdignify.ne* 02/2024 – 04/2025

- Architected and deployed production-ready FastAPI microservices that power high-concurrency AI products, integrating fine-tuned LLM architectures (LLaMA, GPT) to deliver context-aware reasoning and advanced natural language understanding.
- Engineered high-performance Retrieval-Augmented Generation (RAG) systems featuring optimized vector search and automated data ingestion pipelines to drive intelligent discovery, recommendation engines, and automated sales workflows.
- Established comprehensive MLOps and CI/CD frameworks using Docker and automated monitoring to manage the full lifecycle of PyTorch and scikit-learn models, ensuring scalable inference and consistent performance across e-commerce and lead-generation platforms.

Skills & Technologies

Backend & API Development
FastAPI, Flask, PostgreSQL, MongoDB, Redis, Redshift, SQLAlchemy, GraphQL, RESTful APIs

Authentication & Security
OAuth 2.0, Multi-tenant Architecture, OpenID Connect, Secure API Gateway Management, Encrypted ETL Pipelines

MLOps & DevOps
Docker, GitHub Actions, MLflow, Prometheus, Grafana

AI / LLM & Agent Systems
LLM Orchestration (OpenAI, Claude, Groq), Agentic Workflows, RAG Architecture, Vector Databases, Neural Search, LangGraph

Data Engineering & MLOps
Docker, CI/CD (GitHub Actions), MLflow, Prometheus, Grafana, Distributed Inference, Automated Data Ingestion

Projects

AISA Whispar — Real-Time AI Voice Copilot, *Mactay | Alternate Bank* 06/2026 – Present

- Built a real-time voice copilot for call centre agents, combining Deepgram Nova-3 and Google Chirp 3 for live multilingual transcription (English, Yoruba, Hausa) across Twilio, 3CX, and browser channels.
- Implemented a RAG-powered suggestion pipeline surfacing contextual agent cues mid-call using pgvector semantic retrieval and an async worker queue.

AISA Whispa Chat — Dual-Interface Call Centre Support System, *Mactay* 05/2026 – Present

Built an agent-facing chat interface for real-time customer interaction assistance, paired with a dedicated admin control panel for knowledge base management and monitoring.

Omni-Channel AI Support Agent, *Dex, Afriex* 05/2025

- Built a support agent across chat and voice channels implementing a four-stage pipeline (input guardrails → scenario routing → RAG + LangGraph orchestration → output guardrails).
- Automated ticket generation with AI tag classification and DeepEval-powered RAG evaluation for continuous knowledge base monitoring.
- Built a versioned, reproducible evaluation framework with deterministic tool-mocking and multi-turn conversation scoring for regression-safe testing.

Mailman AI — Agentic Email Intelligence System, *Mactay* 03/2026 – Present

- Built a production-grade agentic email system for Gmail and Outlook featuring a multi-format knowledge base (PDF, DOCX, CSV), semantic RAG retrieval over pgvector, and an LLM pipeline classifying intent, urgency, and sentiment with confidence-scored routing.
- Designed on a hexagonal architecture with an async job queue supporting both auto-reply and human-in-the-loop modes.

AISA Omni-Channel Ecosystem(LMS and Mactay's website), *Mactay* 02/2026

Architected a suite of conversational AI and voice agents leveraging RAG and intelligent pause-handling to automate customer engagement across web platforms and LMS website.

Multilingual AI QA Pipeline, *Mactay* 01/2026 – Present

- Developed a pipeline to analyze and score 5,000+ calls for sentiment and responsiveness, integrating cloud storage from OneDrive, Google Drive, and SharePoint.
- Expanded language support to French, Portuguese, Chinese, and Spanish.

SQL Copilot, *Dex, Afriex* 10/2025 – Present

Architected a custom LLM orchestration framework enabling natural language querying with multi-step reasoning across PostgreSQL, MongoDB, Redshift, and Google Sheets.

HealthLang Al - Indigenous Medical Assistant (MCP) 07/2025 – Present

- Built multilingual medical assistant with LLMs and a RAG pipeline (RAG + MCP)
- Designed scalable backend with FastAPI, LangChain, and vector search (FAISS) for symptom analysis.

Hojo news Podcast System 04/2025

- Developed an innovative Text-to-Speech Podcast System with authentic Nigerian accents for a news Web Application, leveraging YarnGPT and FastAPI to convert news content into natural-sounding audio in multiple Nigerian languages (English, Yoruba, Igbo, Hausa).

Vehicle Rentage AI Platform 10/2025 – Present

- Built two microservices using FastAPI:
- Implemented MLOps structure with monitoring (Prometheus), CI/CD, async endpoints, and Docker containers.

VeraCT Scan News Verification System 06/2024 – 07/2024

- Developed an automated news verification system that validates content against 650+ trusted sources, leverages NLP for entity recognition, and utilizes LLAMA- for contextual analysis.

Fingerprint Biometric Attendance System 2024

Final year project

Open Source, Community , Outreach Contributions and Awards

Top 5 Finalist — One With AI Hackathon 2026 2026

Built FleetGuard, an AI-powered fleet intelligence platform for logistics operators to detect fuel theft, route abuse, excessive idling, and operational inefficiencies. Placed 5th overall among participating teams.

AI Backend Engineer — Omdena (Global Open Source AI Collaboration) Log Normalization & Classification Pipeline 08/2025 – 12/2025 Remote

- Built FastAPI microservices for log normalization and classification pipelines
- Integrated ML pipelines using PyTorch and scikit-learn for anomaly detection
- Collaborated with a globally distributed engineering team using GitHub and Slack

Technical Writer — Medium (AI, Cloud, and MLOps topics)

- Authored 4 technical articles on AI systems, LLMs, cloud architecture, and infrastructure optimization, covering topics such as serverless computing, AI agents, transformers, and cloud scalability.

Contributor, Hugging Face Community — Transformer/NLP projects

- Actively involved in Hugging Face community, contributing to transformer-based NLP projects.

Technology & AI Awareness Facilitator — Secondary School Sensitization Programs, *Community Outreach & STEM Education* 2025 Lagos

- Facilitated AI and technology awareness sessions for secondary school students

PLATFORMS & CERTIFICATIONS

Hugging Face 2025
Transformers, Model, Pipelines, Encoder-Decoder Architectures, Huggingface models utilization, Fine-tunning models, Huggingface community, Projects

NVIDIA (Ongoing) Present
Building Real-Time Al Video Applications

DataCamp 2024
NLP Course Completion

Udemy 2023 – 2024
Python Programming

Cisco 2022
Introduction to Cyber Security (Badge Earned)