INIOLUWA OLUFEMI

Under G, Ogbomosho | inioluwaolufemi7@gmail.com | GitHub | LinkedIn

PROFILE SUMMARY

AI/ML and Generative AI developer specializing in Retrieval-Augmented Generation (RAG), NLP, and applied machine learning, currently completing a Computer Science degree at Ladoke Akintola University of Technology (LAUTECH). A graduate of the DeepTech (3MTT) Generative AI program with hands-on experience designing production-style RAG pipelines, self-evaluating AI agents, and deployed ML applications. Skilled at translating raw data and unstructured documents into grounded, traceable, and measurable AI systems.

CERTIFICATION

DeepTech Program (3MTT) *2025*

Graduate, Generative AI — Credential ID: FE/25/2779932023 (verify: app.3mtt.training/verify?id=FE/25/2779932023)

TECHNICAL SKILLS

AI & Machine Learning: Generative AI, LLMs, Transformers, Retrieval-Augmented Generation (RAG), Prompt Engineering, NLP, Semantic Search & Embeddings, Model Evaluation (ROUGE metrics), Structured/JSON Output Design, Data Engineering

Programming: Python, JavaScript, HTML, CSS

AI Tools & Frameworks: TensorFlow, PyTorch, Scikit-learn, Hugging Face Transformers, Sentence-Transformers, FAISS, OpenAI API, Google Gemini API, Gradio, Streamlit, FastAPI

Deployment & Infrastructure: Hugging Face Spaces, Render, Docker, Lovable, API Integration, Model Deployment, Virtual Environments

Version Control: Git & GitHub

Other: Data Cleaning & Preprocessing, Dashboard/Analytics Development

PROJECT EXPERIENCE

Hybrid RAG-Powered PDF Summarization Engine

- Accomplished a document summarization pipeline that condenses long-form PDFs without losing context, measured by built-in ROUGE-1/ROUGE-L, compression ratio, and coverage evaluation surfaced live in the app, by engineering dynamic model routing (lazy-loading between flan-t5-small and flan-t5-base based on document length),

semantic chunking with centroid-based retrieval using sentence-transformer embeddings, and an anti-redundancy filter enforcing a 0.9 similarity threshold, deployed via Gradio (huggingface.co/spaces/oluinioluwa814/PDF).

Self-Improving Telecom Support RAG Agent

- Accomplished a telecom support agent that cites its sources and refuses to hallucinate, measured by a groundedness-scoring "AI Judge" that discards and regenerates unsupported claims and a per-response evidence_used array tracing every diagnosis to specific source documents, by building RAG over real-world telecom ticket data with FAISS, Sentence-Transformers, HNSW indexing, and semantic deduplication, adding issue-category classification (network, billing, SIM) with structured JSON-style outputs, and an autonomous memory loop that embeds high-scoring (9/10+) resolutions back into long-term knowledge (huggingface.co/spaces/oluinioluwa814/Telecommunications; huggingface.co/spaces/oluinioluwa814/Telecom).

Wealth Management Advisor — Finance RAG Assistant

- Accomplished a finance Q&A assistant that grounds its answers in a searchable knowledge base, measured by a live, publicly running Hugging Face Space where each question retrieves relevant information before generating a clear, plain-language response, by building a retrieval-augmented Python backend with a custom HTML/CSS/JavaScript interface, containerized with Docker (huggingface.co/spaces/oluinioluwa814/WealthMangementAdvisor).

Revenue Radar — Retail Sales Analytics Dashboard

- Accomplished a production-ready analytics layer replacing manual spreadsheet reporting, measured by real-time tracking of revenue, order volume, and average order value across an 8,400-row retail transaction dataset, by parsing and persisting raw tab-delimited data into structured storage (rather than re-parsing on every load) and building interactive filtering by date, priority, category, and segment (revenue-radar-gui.lovable.app).

Fake News Detection App

- Accomplished an explainable news-credibility tool that flags articles as likely real, likely fake, or balanced, measured by real-time classification of pasted articles with visible reasoning behind each verdict (not just a label), by training an ML classification model with a text vectorizer and deploying it as an interactive Streamlit web app on Render (fake-news-6tab.onrender.com).

STRENGTHS

- Strong problem-solving mindset
- Data cleaning
- Fast learner and adaptable to new technologies
- Effective communicator
- Passionate about AI research and innovation

EDUCATION

Ladoke Akintola University of Technology (LAUTECH) *Oct 2024 – Aug 2028*

B.Tech, Computer Science — In Progress

Cephas ICT Hub *Aug 2026 – Nov 2026*

Undergraduate Trainee, Machine Learning — applied ML/data projects including Revenue Radar analytics dashboard