

Nduka Nwagbo

AI Engineer • Birmingham, UK

nwagbondukajr@gmail.com

ndukanwagbo.com

github.com/Nduka99

www.linkedin.com/in/nduka-nwagbo-342b16392

huggingface.co/nduka1999

Profile

I turn open-weight models and open-source tools into AI that works in any industry: running where your data lives, grounded in your own documents, and tested before anyone relies on it.

I build AI that professionals can check before they trust it, and I take it through the whole lifecycle: data, modelling, validation, explanation, deployment and handover. I test models against real-world data, not only a test set, and I write down what each system cannot do.

Before AI, I studied communication and multimedia design and worked in IT support, explaining fixes to people who are not technical. I still write for the people who will use the system.

Selected projects

US accident impact risk

Dec 2025 – Oct 2026 • MSc AI coursework, rebuilt independently

Found that the popular “accident severity” label mostly tracked which traffic feed filed the report, then rebuilt an honest model on 3.2 million accidents from one consistent feed, explained it, and made it run entirely in the browser.

Result: 0.852 ROC-AUC on later, unseen months, with no server

us-accident-impact-risk.nduka-nwagbo.workers.dev

github.com/Nduka99/accident_severity_xai

Consumer complaint triage

Mar – Oct 2026 • MSc Deep Learning for NLP coursework

Classifies US consumer financial complaints into ten products with distilled transformers, finds the relevant CFPB examination procedures, and decides whether each case is handled automatically or sent to a person.

Result: Macro-F1 0.757 on 274,065 held-out complaints, up from 0.689

complaint-triage.nduka-nwagbo.workers.dev

github.com/Nduka99/customer_complaint_triage

PetroKnowledge

Oct 2026 • Personal project, independent of NNPC

Answers questions about public NNPC information with local open-weight models. Every claim carries a quote that code checks word for word in its source, figures are checked again by a second model, and unsupported questions get a clear refusal.

Result: Runs on local open-weight models; every quote is checked against its source

github.com/Nduka99/nnpc_ai_assistant

PPE compliance detection

Mar – Oct 2026 • MSc Deep Learning coursework

Merged three public datasets into 10,864 images, trained YOLO11 and served it as an app that flags missing hard hats and vests. An independent real-world benchmark exposed a model that missed real vests; two fixes raised real-world vest recall from 0% to 50.5%.

Result: Real-world vest recall raised from 0% to 50.5% at 80% precision

ppe-detection-blush.vercel.app

github.com/Nduka99/ppe_compliance_detection

Agentic menu development

Jul – Sep 2026 • MSc dissertation with a UK hospitality consultancy (paper mark 85)

My MSc dissertation. A five-stage agentic workflow turns a hospitality client's intake form into a consultant-grade menu brief. Each stage reads only what the approved stage before it hands over. Open-weight models draft, another model family reviews, code checks what they miss, and a consultant approves each step.

Result: Judged usable with minor edits by a practising consultant

SpineLens AI

Jun 2026 • Team entry: West Midlands winners, UKREiF 2026 Next-Gen Challenge

The evidence toolkit behind The Innovation Spine, a team proposal to reconnect Birmingham's Knowledge Quarter to the city centre. Route legibility scoring, wayfinder placement, traffic and collision evidence and a costed plan, shipped as an interactive map.

Result: The route ranking held in 86% of 4,000 random weightings

spinelens-ai.pages.dev

github.com/Nduka99/spinelens-ai

ApexMind

Jun 2026 · BCU AI Hackathon 2026: winning team

A four-hour build for the BCU AI Hackathon 2026. Answer 100 multiple-choice questions using only open models of 8 billion parameters or fewer, grounded in Wikipedia, with confidence scores and repeatable runs.

Result: 100% on a hand-checked 20-question set, up from 90% without retrieval

github.com/tejiri-code/BCU-AI-Hackathon-2026

QHybrid Lab

Jul – Oct 2026 · Independent research

A web service that runs a quantum-kernel classifier against classical models on the same rows, features and tuning folds, for a cybersecurity and a medical dataset, and serves a calibrated network-intrusion triage model. Every comparison can be re-run from the public code.

Result: With a tuned bandwidth the quantum kernel tracked the RBF SVM, at or below the best classical model in 31 of 40 runs

github.com/Nduka99/quantum_hybrid_lab

Experience

Recruitment Ambassador

Birmingham City University · 2025 – present

Explain technical programmes to prospective students and families at open days and campus tours.

IT Support Technician

Bani Law Consulting, Abuja · 2024 – 2025

Kept systems reliable for non-technical staff and explained fixes clearly to users and management.

Education

MSc Artificial Intelligence

Birmingham City University · 2025 – 2026

Distinction (formal award pending). Dissertation paper 85, viva 81. Modules in applied AI, natural language processing, computer vision, advanced machine learning, and AI ethics, bias and governance.

BSc (Hons) Communication and Multimedia Design

American University of Nigeria · 2016 – 2020

Recognition

UKREiiF Next-Gen Challenge, West Midlands winners, 2026. With my team. Our Knowledge Quarter proposal was presented at Birmingham City Council.

BCU AI Hackathon, winners, 2026. With my team.

PGXPO, Silver award, 2026. For my research poster.

Certifications

CompTIA Security+ (SY0-701) and Network+ (N10-009), valid to 2027

Certificate in Cybersecurity, University of Windsor, 2025

Linux Foundation LFS151 Cloud Infrastructure, LFS158 Kubernetes and LFS250 Kubernetes and Cloud Native Essentials, 2025

Skills

Machine learning

PyTorch, Transformers, scikit-learn, LightGBM, XGBoost, CatBoost, YOLO11, knowledge distillation

Evaluation and explainability

SHAP, Integrated Gradients, time-based splits, bootstrap intervals, McNemar's test, calibration, independent benchmarks, evaluation harnesses, pre-registered evaluation

LLMs and retrieval

local open models with llama.cpp, GGUF quantisation, agentic workflows with human approval, LLM review across model families, RAG, BM25, pgvector, cross-encoder reranking, schema-constrained generation, document extraction

Software and deployment

Python, TypeScript, SQL, FastAPI, React, Postgres, Docker, GitHub Actions, Cloudflare, Hugging Face Spaces, AWS, Terraform, Playwright, pytest