

Timothy Ojebiyi

7 Homestead Way, Croydon, CR0 0BG | +44 78732 4900 |
tobiojebiyi@gmail.com | [LinkedIn Profile](#) | [GitHub](#) | [Website](#)

Objective

Computer Science graduate specialising in Artificial Intelligence and Full-stack Development. Builds AI-powered products end-to-end from agentic pipelines and LLM tool-use to production infrastructure with a track record of solo ownership across the full stack, from database design through deployment.

Education

BSC(HONS) COMPUTER SCIENCE | 2022 - 2025 | THE UNIVERSITY OF THE WEST OF ENGLAND, BRISTOL
Grade: 2:1 (Honours)

Experience

FULL-STACK ENGINEER INTERN | CARBONTRAC | JULY 2025-SEPTEMBER 2025

- Developed a PostgreSQL database with a 10-table star schema storing product emissions data, supplier sources, delivery routes, and user shopping histories.
- Designed and built the nutrition tracking ML model that analyses 46 nutritional parameters across macronutrients, micronutrients, and functional health tags (diabetes-friendly, pregnancy-supportive).
- Built a personalised swap recommendation engine that detects nutrient gaps from shopping history and suggests alternatives e.g., flagging selenium deficiency.

UX & ACCESSIBILITY INTERN | NOMENSA | JULY 2024-AUGUST 2024

- Led UX research for an AI-powered healthcare product, conducting stakeholder workshops and user interviews that shaped the product's core user flows.
- Developed user personas, service blueprints, and journey maps that identified 4 critical pain points in the patient experience.
- Designed and prototyped interface solutions in Figma, iterating based on usability testing feedback.

Projects

ROOTLAYER - AI-POWERED WEB COMPLIANCE AUDITOR

Built an AI-powered compliance auditing platform that helps teams find and remediate accessibility (WCAG 2.2 AA) and privacy (GDPR/CCPA/TCF) issues. Async workers run Playwright-based scans combining axe-core, custom checks, and LLM analysis to generate structured violations with severity ratings, remediation guidance, and shareable reports, with support for multi-tenancy and white-label branding for agencies.
Technologies: Next.js, FastAPI, Celery, Redis, PostgreSQL, Playwright, Anthropic SDK, OpenAI SDK

CONTROLNET IMPLEMENTATION - PARITY-VERIFIED

A from-scratch PyTorch reimplementation of ControlNet, validated by testing that its outputs match the Hugging Face diffusers reference implementation to $\sim 1e-4$ proving mathematical correctness rather than treating it as a black box. Trained on 45,000 COCO images, improving edge fidelity (F1 0.34 $\rightarrow$ 0.79) without degrading base model text.
Technologies: Python, PyTorch, Hugging Face Diffusers, Accelerate, pytest.

UNSAID - SEMANTIC 10-K DISCLOSURE-REMOVAL DETECTOR.

Built a pipeline that detects when public companies quietly remove or soften risk disclosures between consecutive annual SEC filings a pattern academic research shows is heavily underpriced by markets. The system uses local semantic embeddings to align disclosure units year-over-year, then uses Claude to classify each as removed, softened, reworded, retained, or new. Validated against the Silicon Valley Bank case study, correctly flagging the removal of interest-rate hedging disclosures from a filing submitted 14 days before its collapse.
Technologies: : Python, FastAPI, Anthropic SDK (Sonnet + Opus), sentence-transformers, Next.js

Skills & Abilities

Languages: Python, TypeScript, JavaScript, C++, SQL
Backend: Django, FastAPI, Node.js, Express, REST APIs, PostgreSQL, Redis, Celery
Frontend: React, Next.js, Tailwind CSS, Responsive Design
ML/AI: PyTorch, TensorFlow, CNNs, Transfer Learning, Data Analysis (pandas, NumPy)
Tools: Docker, Git, AWS, Azure DevOps, Postman, Figma