

Olajumoke Taiwo, Backend AI Engineer

Ontario, Canada, +1(437)-435-0371, taiwoolajumoke24@gmail.com

PROFILE

Data Scientist and AI Engineer with 5+ years developing predictive models and machine learning solutions that drive measurable business impact. Expert in deep learning, time-series forecasting, statistical modelling, building end-to-end ML pipelines from data collection to model deployment using TensorFlow, PyTorch, and cloud-based ML infrastructure. Combines technical depth in neural networks and transformer architectures with business acumen to solve complex problems across environmental analytics, healthcare insights, and financial modelling.

TECHNICAL SKILLS

Python (Pandas, NumPy, Matplotlib, Seaborn), R, SQL, Statistical Analysis	Expert	Supervised Learning, Deep Learning, Time-Series Forecasting, Neural Networks, Ensemble Methods	Experienced
TensorFlow, PyTorch, Scikit-learn, Keras, XGBoost, LightGBM, Prophet, ARIMA	Expert	Pinecone, ChromaDB, FAISS, Similarity Search, Semantic Retrieval, Optuna	Experienced
BERT, Transformers, LSTM, RNN, LangChain, OpenAI API, Ollama, RAG Systems, Vector Embeddings, LiteLLM	Expert	AWS (SageMaker, Lambda, S3), MLflow, Docker, Model Deployment, A/B Testing	Experienced
Power BI, Tableau, Plotly, Interactive Dashboards, Statistical Reporting	Expert	Apache Airflow, ETL Pipelines, Data Validation, Feature Engineering, Model Monitoring	Skillful

EMPLOYMENT HISTORY

Jan 2026 — Present

Backend AI Engineer, Hello Biome

- Architected scalable backend infrastructure powering LLM-based AI solutions for microbiome research, ensuring high availability and production-grade performance across AWS serverless architecture
- Developed production AI interpretation pipeline processing microbiome data through AWS Lambda and multiple LLMs, generating personalised skin health insights and evidence-based recommendations
- Implemented retrieval-augmented generation using pgvector and vector embeddings to ground AI outputs in 500+ scientific papers, enhancing interpretation accuracy while reducing hallucinations through vector similarity search
- Led model evaluation initiative across 35+ language models, establishing testing frameworks and training team members on LLM assessment methodologies to optimise production deployment and maintain quality standards

Jul 2025 — Dec 2025

Independent AI Engineer

- Designed and developed AI-powered applications including DendriteAI (Knowledge management and meeting intelligence system using RAG pipelines) and Closer (Faith-based support and wellness app)
- Built and experimented with agentic AI workflows, enabling multi-step reasoning, tool use, and autonomous task execution within LLM-based systems
- Implemented retrieval-augmented generation (RAG), embeddings, and vector databases to power intelligent search and response systems

Mar 2025 — Jul 2025

Fullstack Data Scientist, Panthalassa

- Enabled proactive environmental compliance by building emission quantification models for NOx, SOx, PM2.5, and CO2 from vessel operations, creating pollution hotspot maps that helped ports identify critical intervention areas and optimise shore power infrastructure
- Reduced carbon footprint assessment time from weeks to hours by developing automated calculation and forecasting systems that process client environmental data with scenario-based models, enabling data-driven carbon reduction planning across different intervention timelines
- Improved port air quality management by creating time-series forecasting models that predict port-level air quality using historical vessel movements, weather patterns, and emission measurements, enabling proactive mitigation strategies

- Strengthened the climate risk assessment service by developing scenario-based infrastructure risk models using climate datasets with uncertainty quantification, supporting strategic planning decisions

Jan 2023 — Jan 2025

Machine Learning Engineer, Reagle Consulting

- Reduced fraud losses by 65% by building ensemble-based fraud detection models (XGBoost + Logistic Regression + Neural Networks) achieving 84% accuracy, enabling real-time transaction screening and reducing false positives by 34%
- Improved inventory planning accuracy by 12% through bidirectional LSTM networks with attention mechanisms for demand forecasting models that reduced stockouts and overstock situations
- Enhanced customer service efficiency by 23% by fine-tuning BERT models for automated sentiment analysis of support tickets, enabling priority routing and reducing response times
- Increased customer retention by 15% through gradient boosting churn prediction models achieving 0.89 AUC, with SHAP-based insights that guided targeted retention campaigns
- Maintained 99.2% model uptime by deploying comprehensive model monitoring including data drift detection, performance tracking, and automated alerting systems for production models

Feb 2021 — Nov 2022

Data Scientist, Notyl Services

- Optimised vaccine distribution efficiency by 18% through demand forecasting models that improved resource allocation across 4 countries, ensuring better coverage in underserved areas
- Increased program resource allocation efficiency by 22% using Random Forest and XGBoost algorithms to identify optimal intervention targets, maximizing impact per dollar spent
- Streamlined data operations by creating Apache Airflow ETL pipelines that automated monthly survey processing, reducing manual effort by 75% and improving data quality consistency
- Improved program evaluation rigor by implementing A/B testing frameworks with proper statistical controls, enabling evidence-based decision making for demographic-specific interventions

EDUCATION

Sept 2022 — Sept 2023

M.Sc. Artificial Intelligence and Data Science, Keele University

Distinction

Sept 2017 — Jun 2021

B.Sc. Public Health (Biostatistics), Babcock University

PROJECTS

Closer (Christian Wellness Mobile App)

Tech Stack: Python, RAG, Pinecone, OpenAI API, Transformer Models, Vector Embeddings

- Processed extensive theological knowledge by implementing vector embedding systems that enable semantic search across comprehensive Christian therapy datasets
- Delivered personalised faith-based guidance via emotion-based recommendation engines that analyse user input for contextually appropriate responses
- Impact: Provided 24/7 spiritual support through RAG-powered AI chatbot processing comprehensive Christian therapy datasets, reducing barriers to accessing faith-based wellness support

Mindwave (Psychological First Aid Platform)

Tech Stack: BERT, GPT-3.5, NLP, Transformers, Sentiment Analysis, PyTorch

- Fine-tuned BERT model on conversational dataset (50,000+ mental health dialogues) for emotional distress detection achieving 78% accuracy in 8-class emotion classification with stratified sampling
- Provided accessible mental health support through transformer-based chatbots trained on professional therapy sessions, offering immediate psychological first aid
- Enhanced mood pattern recognition via sentiment analysis pipelines that process journal entries, helping users and clinicians track mental health trends over time using time-series analysis and trend detection
- Impact: Improved mental health crisis response by developing BERT-based emotional distress detection with 78% accuracy, enabling early intervention and support for at-risk individuals

SkinIQ (Skin Analysis Platform)

Tech Stack: TensorFlow, CNN, Computer Vision, MongoDB, Python, Image Processing

- Implemented image pre-processing pipeline with rotation, scaling, and colour normalisation techniques increasing model robustness by 30% on validation dataset
- Enabled comprehensive skin analysis through computer vision-based feature extraction systems analysing texture, colour, and pattern characteristics
- Impact: Simplified dermatological screening by building a CNN achieving 89% accuracy across 27 skin conditions, helping users identify potential skin issues and get appropriate care recommendations

RESEARCH AND PUBLICATIONS

Emotion-Aware Psychological First Aid: Integrating BERT-based Emotional Distress Detection with PFA-GPT Chatbot for Mental Health Support (2025)

Co-author: Dr. Baidaa Al-Bander

Cognitive Computation and Systems Journal

Novel approach combining emotional distress detection with automated psychological support using deep learning models

COURSES

AI Engineering, IBM