

Godswill Ikwan

+1 (203) 503 7627 | godswillikwan01@gmail.com | [linkedin.com/in/godswill-ikwan](https://www.linkedin.com/in/godswill-ikwan) | github.com/awesomeguss | datascienceportfol.io/godswillikwan01

EDUCATION

University of New Haven

M.S. in Data Science, GPA: 3.89

West Haven, CT

December 2024

Covenant University

B.Eng. in Information and Communication Engineering

Ota, Nigeria

December 2020

EXPERIENCE

Token Metrics

September 2025 – March 2026

Machine Learning Engineer

West Haven, CT

- Built distributed ML forecasting and real-time trading pipelines using Python, XGBoost, and PostgreSQL, achieving 73.3% directional accuracy across 98,133 forecasts.
- Provisioned scalable AWS compute infrastructure using Terraform to manage containerized time-series forecasting models and high-throughput DeFi data ingestion workloads on Kubernetes (EKS)
- Architected robust CI/CD pipelines via GitHub Actions that executed automated pytest suites to validate Hyperliquid SDK protocol behavior and applied Alembic schema migrations, reducing model deployment cycles from days to under 20 minutes
- Deployed a production-grade, containerized model serving API utilizing FastAPI and Pydantic, fed by a self-healing WebSocket ingestor processing 5,000+ messages/second and achieving sub-50ms inference latency for 1-minute OHLCV candles.
- Implemented comprehensive ML system observability using Prometheus and Grafana for cluster utilization, alongside Evidently AI to monitor prediction drift, automating retraining triggers and reducing mean time to detection (MTTD) for model degradation by 70%.
- Developed a FastAPI trading service using the Hyperliquid Python SDK and Pydantic validation to expose 6 endpoints for order lifecycle management.

Extern x Pfizer February 2026 – May 2026

AI Engineer Fellow, Pfizer Supply Chain

Remote, West Haven, CT

- Built and optimized an end-to-end RAG chatbot using Python, LlamaIndex, and Docker to cut pharmaceutical document processing latency by 87%.
- Optimized the document pipeline using a rule-based pre-filter and parallel Gemini classification to reduce latency from 60s to 8s.
- Containerized the retrieval system using Docker with a CUDA-enabled Dockerfile and unit tests to ship a reusable open-source Python package.

University of New Haven

January 2025 – January 2026

Research Assistant, Machine Learning & Deep Learning

West Haven, CT

- Developed a regularization framework using the Choquet integral to resolve limited-training-data bottlenecks, boosting classification accuracy by ~5pp across 7 tasks.
- Co-authored a research paper using optimized model aggregation techniques to improve predictive accuracy by 5%, presented at IEEE WCCI 2026.

PROJECTS

Agentic AI Data Warehouse Explorer

November 2025 – January 2026

Python, FastAPI, Snowflake, LangChain, LangGraph

- Built an AI analytics app using FastAPI, Snowflake, and local LLMs to translate natural-language business questions into validated SQL queries.
- Implemented LangGraph planning/execution/repair orchestration using reliability checks to sustain under 5s processing for 10k-row result sets.

TECHNICAL SKILLS

Cloud & Orchestration: AWS (EKS, SageMaker, S3, Lambda, Bedrock), Kubernetes, Docker

Languages & Scripting: Python, SQL

Databases & Data Systems: PostgreSQL, MongoDB, Snowflake

DevOps & Tools: Git, GitHub Actions, CI/CD, FastAPI, MLflow

ML/AI Frameworks: PyTorch, TensorFlow, Scikit-learn, XGBoost, LangChain, LangGraph, Hugging Face

CERTIFICATIONS & ACHIEVEMENTS

AWS Certified Gen AI Developer – Professional (In Progress) | Published Research: Unified Regularization Framework for the Choquet Integral | Dean's List, Master's Program (GPA: 3.89)