ONYEOCHA KESIDE OLATUNBOSUN

APPLIED AI & FORWARD DEPLOYMENT ENGINEER | AI AGENTS, PYTHON & ENTERPRISE INTEGRATIONS

Lagos, Nigeria | +234 814 190 2806 | onyeocha.keside@gmail.com
kesideonyeocha.com | linkedin.com/in/keside-onyeocha-bb48b9298 | github.com/Onyeocha-Keside

PROFESSIONAL SUMMARY

AI/ML Systems Engineer with 3+ years of experience building production ML and intelligent systems end to end - from automated data acquisition and ML-ready pipelines through inference, APIs, cloud orchestration and secure delivery. Built high-volume work-intelligence platforms, AWS-native data and model workflows, production RAG/agentic systems and resource-efficient edge AI. Strong Python and distributed-systems engineering with hands-on ownership across architecture, deployment, observability, reliability and technical leadership.

TECHNICAL EXPERTISE

AI/ML: PyTorch, TensorFlow, Scikit-learn, XGBoost, Hugging Face Transformers, feature engineering, time-series analytics, computer vision, model evaluation, statistical testing

LLM & GenAI Systems: RAG, agentic workflows, LangChain, LangGraph, dense/sparse retrieval, Pinecone, Qdrant, structured outputs, prompt/retrieval evaluation, OpenAI, Anthropic, ElevenLabs

ML Systems & Data: model serving, MLOps, batch data pipelines, training-data preparation, Parquet, DuckDB, PostgreSQL, Redis, SQLAlchemy, Alembic, ONNX, TensorRT, CUDA

Software Engineering: Python, C++, FastAPI, Flask, REST APIs, asynchronous services, SQLite, Vue.js, automated testing, distributed systems

Cloud & Platform: AWS (ECS, Lambda, Step Functions, EventBridge, RDS, S3, ECR, CloudWatch, Secrets Manager, IAM, SSM), GCP (Vertex AI, BigQuery), Docker, Terraform, GitHub Actions, CI/CD, Airflow

PROFESSIONAL EXPERIENCE

Doheney Services | Remote | Part-time **Sep 2025 - Present**

AI/ML & Software Systems Engineer

- Architected and delivered **FocusTime**, a privacy-first, multi-tenant work-intelligence platform that transforms 20M+ high-frequency activity records into explainable deep-work, context-switching, workload, energy and team insights using FastAPI, PostgreSQL, Redis and Vue.js.
- Built an offline-first Python desktop agent for macOS and Windows with OS-level activity/idle detection, encrypted credentials, persistent SQLite queues, fault-tolerant batch ingestion, consent controls and automatic updates.
- Developed a grounded AI layer combining auditable SQL feature engineering, behavioral time-series analytics and rule-based energy modeling with Claude-powered structured narration; enforced schema validation, deterministic fallbacks and strict metric grounding.
- Implemented an ML-ready lakehouse path with transactional outbox semantics, idempotent ZSTD Parquet exports and recovery controls; productionized the platform on AWS with Terraform/GitHub Actions OIDC and nearly 800 automated tests.

Greyzone Limited | Remote | Contract **Dec 2025 - Aug 2026**

Machine Learning Pipeline Engineer

- Engineered automated daily crypto-market intelligence pipelines ingesting news, narratives, whale activity and trend signals from Google Trends, CoinGecko and external APIs; performed preprocessing/EDA and persisted curated ML-ready datasets to S3 and Parquet.
- Built scheduled AWS orchestration for upstream/downstream pipeline dependencies using Lambda, Step Functions and EventBridge, with ECR, Secrets Manager, RDS/S3 integration, Terraform, GitHub Actions and CloudWatch/SNS observability.
- Supported model-training and champion-model workflows by producing reproducible training artifacts, reliable feature/data handoffs and pipeline integrations for models hosted on AWS.
- Designed an automated multimodal tiktok content pipeline that converted aggregated market news and narratives into LLM-generated scripts/dialogue, ElevenLabs voiceovers, stock-media assets and rendered short-form videos, with validation and failure-tolerant persistence and posted directly on tiktok.

PROFESSIONAL EXPERIENCE (CONTINUED)

JeGO | USA | Remote **Nov 2024 - Oct 2025**

AI/ML Systems Engineer

- Architected a unified three-phase agentic RAG platform spanning customer support, EV booking and analytics, consolidating model infrastructure and reducing the production deployment surface.
- Delivered a customer-service inference API with EV-booking integration; optimized caching and retrieval to sustain sub-second responses under concurrent load.
- Engineered applied AI systems spanning low-latency SLAM/multi-sensor navigation and predictive optimization, including solutions that improved solar-tracking efficiency by 34.5% and reduced average delivery time by 22%.

Sanctum | Nigeria **Apr 2024 - Oct 2024**

AI/ML Engineer - Retrieval & Agentic Systems

- Designed and deployed a hybrid dense-sparse RAG pipeline with Pinecone, improving retrieval relevance by 40% across production knowledge bases.
- Built multi-agent workflow infrastructure that reduced manual processing time by 60% through coordinated task decomposition, agent communication and execution.
- Evaluated retrieval improvements across 5,000+ queries using paired t-tests and bootstrapped confidence intervals, and led 10+ technical sessions on RAG and multi-agent deployment for 15+ engineers.

NITDA IT Hub | Lagos, Nigeria **May 2023 - Aug 2023**

Machine Learning Engineer - Edge AI

- Optimized TensorFlow models through quantization and pruning, reducing model size and hardware requirements by 70% while preserving accuracy on resource-constrained edge devices.
- Built automated data/model validation pipelines and developed a native-language keyword-spotting model achieving 95% precision using augmentation and active learning.

RESEARCH

CampusTri-Nav - LiDAR-Free Autonomous Electric Tricycle **Aug 2026 - Present**

Independent AI & Robotics Researcher | Work-in-Progress Manuscript

- Leading systems architecture and evaluation for a low-speed autonomous research tricycle operating within a surveyed University of Lagos campus domain; designed stereo visual-inertial-wheel-GNSS-map fusion and a compact perception stack.
- Authored a safety-gated experimental protocol covering route-grouped splits, leakage prevention, sensor-denial testing, ATE/RPE, mAP, IoU, HOTA/IDF1, calibration, tail latency, power and bootstrap confidence intervals.

LEADERSHIP & COMMUNITY

Cave.AI Hub, University of Nigeria **Jan 2024 - Jul 2025**

Founder & AI Engineering Community Lead

- Founded and led a 50+ member AI engineering community; designed production-ML learning programs, organized monthly RAG/LLM sessions for 40+ practitioners and mentored engineers through hackathons and AI competitions.

EDUCATION & PROFESSIONAL DEVELOPMENT

University of Nigeria, Nsukka | B.Eng., Electronics & Computer Engineering

Fellowship AI, San Francisco - Fellow, Advanced AI Research Program (Jan-Apr 2023) | **Microsoft Responsible AI** - Hands-on Session (Jun 2023)