Emmanuel Osanebi

AI Engineer | LLM Agents, Automation & Production AI (Python, FastAPI)

emmanuelosanebi1@gmail.com | +234 701 147 2548 | LinkedIn: linkedin.com/in/emmanuel-osanebi-1a7042222
GitHub: github.com/emms204 | Portfolio: emms204.github.io | Remote (UTC+1)

SUMMARY

AI Engineer with 5+ years of experience building production LLM agents, chatbots and automation in Python and FastAPI. At MasteryHive AI, built the fraud detection, AML/KYC/KYB, agentic banking reconciliation and reinsurance data products. Previously shipped agentic data pipelines, document-AI workflows and RAG systems, and published a PII-masking SDK on PyPI.

EXPERIENCE

AI Engineer · MasteryHive AI — Aug 2026 to Present
Lagos, Nigeria / Remote

- Built the transaction fraud detection backend end to end, including a streaming service for real-time transaction monitoring.
- Designed and shipped the AML, KYC and KYB backend for identity, customer and business verification and compliance checks.
- Built the banking reconciliation platform: FastAPI microservices on GCP that match bank statements to general-ledger transactions through fuzzy, two-stage rule-based and LLM-assisted matching, with Pub/Sub-fed parallel workers and a dead-letter replay path.
- Designed and shipped ReconPilot, an agentic reconciliation copilot (FastAPI, React/TypeScript, Gemini): the LLM proposes schema mappings, explains evidence and orchestrates approved tools, while all financial calculations stay deterministic in Python, with human approval of every mapping and resolution, an audit trail, a live Server-Sent Events tool timeline, versioned reruns and CSV/JSON/HTML reports.
- Built ingestion, cleaning and validation pipelines for reinsurance bordereaux and treaty data.

AI Engineer · VoterIQ — Mar 2026 to Jul 2026
Remote (India-based team)

- Built an end-to-end political-intelligence pipeline (article ingestion, enrichment, social probing, ranking, database write-back) producing 50 to 100 evidence-scored topics per run from 12 source feeds.
- Engineered durable execution with a Supabase-backed queue, idempotency and a seven-state job lifecycle (queued, claimed, running, completed, failed, deferred, dead-lettered).
- Hardened run reliability with tiered fallback and cooldown logic, including two-level fallback paths for YouTube and X data.
- Automated weekly trend extraction and reranking in a closed V1/V2 loop, cutting manual analyst review by 40%.

AI Software Engineer · Enzagroup Global — Jan 2024 to Dec 2025
Cairo, Egypt / Remote

- Built merchant-onboarding AI that extracts structured data from 20+ document types at 95% accuracy, reducing manual compliance review by 60%.
- Designed agentic document and decision workflows with CrewAI and LangGraph.
- Integrated AI services with identity verification, CRM and internal microservices using FastAPI, async services and caching.
- Deployed AI workloads on AWS (EKS, EC2, RDS) at 99.9% service availability, with monitoring and retraining triggers.
- Added behavioural features and anomaly checks for risk scoring.

AI Engineer · Cali — Jul 2022 to Dec 2023

Hannover, Germany / Remote

- Built LLM backends and multi-agent workflows for large-scale text analysis.
- Delivered a RAG pipeline with LlamaIndex and OpenAI that made document analysis 30% more efficient.
- Improved agent performance by 20% with multi-agent planning and memory, served through Python and Node.js inference backends.

Machine Learning Engineer · Omdena — Jul 2021 to Dec 2022

Global / Remote

- Developed a computer-vision model for weld-defect detection and classification that improved accuracy by 20%, plus APIs that cut training and retraining time by 10%.
- Collected and annotated data for multi-label, attention-based classification of GESI content.

SELECTED PROJECTS

DevGenie: Multi-Agent Full-Stack App Generator — Aug 2025 to Nov 2025

Agents generate React/Next.js frontends and FastAPI backends in parallel from a single prompt, using typed structured outputs.

Python, FastAPI, TypeScript, Next.js · github.com/emms204/DevGenie

AnotiAI PII Masker: Python SDK on PyPI — Nov 2025 to Mar 2026

Hybrid transformer, QA and rules-based PII detector with reversible masking (mask and unmask via a PII map), shipped as the *anotiai-pii-masker* SDK with cloud mode and local fallback, plus an API and web platform.

Python, Hugging Face Transformers, FastAPI, RunPod, GCP · pypi.org/project/anotiai-pii-masker · platform.anotiai.com

Grunsco-Automate: Voice-of-Customer Automation

n8n and FastAPI system that classifies customer reviews with OpenAI and Pydantic structured outputs plus regex guardrails, drafts replies for Slack approval, escalates sensitive cases and mines batch insights.

n8n, FastAPI, PostgreSQL, OpenAI, Pydantic, Slack · github.com/emms204/Grunsco-Automate

WhoseDEV: Privacy-First AI Code Editor

Local-first AI coding agent with sensitive-code masking, multi-model routing and bring-your-own API keys.

whosedev.com

ScriptGen: LLM Script-Generation Framework

Modular LLM pipeline that turns loglines into scenes while maintaining character, plot and setting "bibles" for consistency.

Python, LLMs · github.com/emms204/ScriptGen

SKILLS

LLMs & Agents: LLMs, prompt engineering, RAG, multi-agent systems, tool-use agents, function calling, structured output and JSON/schema-constrained generation; LangGraph, LangChain, CrewAI, LlamaIndex, SmolAgents; OpenAI, Gemini, Anthropic

Languages: Python, SQL, TypeScript, JavaScript, Node.js, Bash, C++

Backend, APIs & Automation: FastAPI, Express.js, REST and WebSocket APIs, webhooks, auth, rate limiting, PostgreSQL, Supabase, n8n

Infrastructure & MLOps: Docker, Kubernetes, AWS (EKS, EC2, RDS, Amplify, Cognito, AppSync), GCP, RunPod, Airflow, Kubeflow, CI/CD, model serving, observability

Applied ML: PyTorch, TensorFlow, Hugging Face Transformers, scikit-learn, XGBoost, LightGBM; classification, anomaly detection, ranking, segmentation, calibration, time-series validation

LLM Training & Evaluation: SFT, instruction-data preparation, DPO/RLHF principles, PR-AUC/ROC-AUC, threshold calibration, latency benchmarking, drift monitoring, regression and adversarial testing

Ways of Working: async remote collaboration, code reviews, agile, startup environments

EDUCATION

BSc Computer Science · University of Ibadan, Oyo, Nigeria — 2025