

SOLOMON ADENUGA

AI Engineer • Machine Learning • Full-Stack Python Developer

Lagos, Nigeria | solomonadenuga8@gmail.com | +234 802 371 0562

GitHub: github.com/TheLogeek | LinkedIn: [solomon-adenuga-6251a5316](https://www.linkedin.com/in/solomon-adenuga-6251a5316) | Portfolio: solomonadenuga.cv

PROFESSIONAL SUMMARY

Results-oriented Machine Learning Engineer and Data Scientist with hands-on experience designing, building, and deploying end-to-end ML pipelines, LLM-powered applications, RAG systems, and AI-driven web platforms. Proficient in Python, Scikit-learn, TensorFlow, FastAPI, and Streamlit with a track record of translating complex datasets into actionable business insights. Founder of LogeekMind — a free, full-stack AI academic productivity platform with 11 specialized tools, including MindMate, an AI chatbot built to support university students — serving 120+ active users — and Scrylo, a commercialized B2B lead generation engine. Adept at LLM orchestration, prompt chain engineering, and multi-tier GenAI architectures. Actively seeking remote, freelance, contract, or internship roles in machine learning engineering, data science, or AI-driven product development.

SKILLS

- **Programming:** Python, NumPy, Pandas, Scikit-learn, TensorFlow, PyTorch, NLTK
- **ML & Data Science:** Predictive Modeling, Classification, NLP, RFM Analysis, Time-Series Forecasting, Feature Engineering
- **GenAI & LLM Engineering:** LLM Orchestration, RAG Pipelines, Multi-Prompt Chains, Context Engineering, Function Calling, Structured Outputs, OpenAI / Claude / Groq / Ollama APIs
- **Visualization:** Matplotlib, Plotly, Streamlit Dashboards, Interactive Data Visualization
- **Deployment & Tools:** Streamlit, FastAPI, REST APIs, Telegram Bot API, Supabase, Git, GitHub
- **Analytics:** Data Cleaning, EDA, Business Intelligence, Automated Reporting, Pipeline Development

WORK HISTORY

Founder & Independent ML Engineer / Data Scientist

Jan 2023 – Present

Independent Projects — Lagos, Nigeria (Remote)

- Engineered FORZA AI and NBA Prophet Pro, multi-league sports analytics engines spanning the EPL, La Liga, Bundesliga, and Serie A plus 5 seasons of NBA data; built a 200+ variable feature matrix (fatigue mapping, dynamic Elo rating) and multi-target stacked ensemble models (XGBoost, LightGBM, Ridge) with auto-invalidating caching to protect API limits.
- Developed and commercialized Scrylo, a private B2B lead intelligence application launched under a \$99 lifetime license model to save early-stage startups from recurring SaaS database fees.
- Built DriftShield AI, a production model-monitoring engine pairing an asynchronous FastAPI background worker with Benjamini-Hochberg corrected KS/Chi-Square drift tests, validated via a custom offline harness achieving a <5% false-positive rate.
- Engineered LogeekMind, a free full-stack AI academic platform with timed exam simulators and notes-to-audio converters; scaled to 120+ registered university students with ~50% daily active engagement.
- Built and scaled a Smart Expense Tracker Telegram bot using NLP-driven parsing and low-latency persistent storage to automate mobile transactional accounting for active users.

Founder & Independent Web Developer

Jan 2022 – Present

Independent Projects — Lagos, Nigeria (Remote)

- Delivered 20+ client web development projects, including 2 e-commerce websites, multiple portfolio sites, a progressive web app (PWA), and product landing pages — meeting all deadlines and specifications.
- Translated client business requirements into functional technical deliverables, maintaining clear communication throughout each project lifecycle.

Data Science Intern — Codveda Technologies

Oct 2025 – Dec 2025

Remote

- Automated end-to-end data cleaning, EDA, and BI reporting workflows using Pandas and NumPy, delivering interactive Streamlit dashboards that surfaced actionable insights for business stakeholders.
- Designed and optimized classification and regression models, applying cross-validation and hyperparameter tuning to improve model performance and support data-driven decisions.

Machine Learning Intern — CodeAlpha

Jan 2026 – Present

Remote

- Built and deployed end-to-end ML pipelines and time-series forecasting models using Scikit-learn and TensorFlow, covering ingestion, preprocessing, training, evaluation, and production deployment.
- Implemented advanced feature engineering and structured data cleaning — missing-value handling, categorical encoding, domain-specific interaction features — measurably improving baseline model precision and accuracy.

FEATURED PROJECTS

Scrylo — Local AI Lead Intelligence & B2B Sales Engine

[Live Site: scrylo.site](https://scrylo.site)

Python, SQLite, EasyOCR, LLM Orchestration, Jinja2, CLI Engineering

- Engineered a local-first B2B sales application enabling users to input target queries (e.g., “SaaS CEO in London”) to scrape multi-source leads on client-side hardware with zero external server dependencies.
- Built a pipeline featuring automated MX-record domain verification, an NLP-driven Ideal Customer Profile (ICP) grading system, and an EasyOCR parser to ingest lead data from business cards into relational SQL rows.
- Implemented a four-tier fallback LLM generation architecture (Groq Llama 3.3 / Ollama / Jinja2) that synthesizes hyper-personalized cold pitches and auto-delivers them via randomized SMTP channels.

LogeekMind — AI-Driven Academic Productivity Platform

[Live Site: logeekmind.xyz](https://logeekmind.xyz)

Next.js PWA, Python, FastAPI, Supabase, Google Gemini API, Groq Cloud API, LLM Orchestration

- Decoupled a high-velocity Next.js PWA frontend from a localized Python/FastAPI REST API backend to support real-time data streaming across an automated Notes-to-Audio converter and a persistent Socratic Tutoring interface.
- Designed structured output pipelines powering a Quiz Generator and Exam Simulator with countdown timers, persistent session management, and auto-submit logic via the Groq cloud API and strict context engineering.
- Deployed live to 120+ registered university students to automate exam simulation, homework analysis, and textbook-to-audio study guide generation; distributed free with an Android wrapper for mobile access.

RAG-Powered Proposal Synthesizer — AI Freelance Proposal Engine

[GitHub Repo](#) | [Live App: rag-powered-proposal-synthesizer.vercel.app](https://rag-powered-proposal-synthesizer.vercel.app)

Python, React 18, Vite, Vercel Serverless, Groq (Llama 3.1), NumPy, @xenova/transformers (WASM), SSE Streaming

- Designed and deployed a full-stack RAG system that embeds job descriptions in-browser via a quantized WASM model (all-MiniLM-L6-v2), performs cosine similarity retrieval over a pre-built embeddings index using NumPy, and streams a grounded proposal token-by-token via Server-Sent Events — eliminating cloud embedding infrastructure.
- Implemented heading-boundary chunking over markdown knowledge base files; top-4 retrieved chunks feed a structured Groq (llama-3.1-8b-instant) prompt chain with relevance scores emitted as a leading SSE event for real-time traceability.
- Decoupled a React 18 + Vite frontend from a Python serverless backend via a 3-panel dashboard (KB Browser, Job Input, streaming Proposal Output); the WASM model caches in IndexedDB for sub-second repeat-query latency.

Edge Proctoring & Attention Tracking Engine

[GitHub Repo](#) | [Live App: edge-proctoring-system.vercel.app](https://edge-proctoring-system.vercel.app)

Python, React 18, Vite 5, Pyodide (WebAssembly), MediaPipe FaceMesh, Scikit-learn, NumPy, Vercel

- Built a fully browser-based proctoring engine where 478 facial landmarks from MediaPipe FaceMesh are piped into a Python runtime executing in WebAssembly via Pyodide, performing all matrix math and behavioral classification on-device with zero server costs and zero video data leaving the client.

- Offline training pipeline synthesizes 5,000 head-pose samples, engineers 30 landmark-derived features, and compares Logistic Regression vs. Random Forest via GridSearchCV; weights export as JSON for pure-Python inference inside Pyodide with no sklearn dependency at runtime.
- Hybrid ensemble: ML classifier takes precedence above 65% confidence, falling back to a rule-based yaw/pitch/gaze engine with a 24-frame violation buffer to debounce jitter. Deployed on Vercel with a <1 MB bundle and real-time canvas overlay.

DriftShield AI — Validated Production Model Monitoring Engine

GitHub Repo

Python, FastAPI, SQLite, Streamlit, SciPy, NumPy, Pydantic, Pytest

- Architected a lightweight, local-first model monitoring engine that runs next to live inference pipelines, tracking data distributions and model health with \$0 cloud infrastructure overhead.
- Built an asynchronous FastAPI background worker to ingest payloads into a structured SQLite ledger without blocking inference latency. Applied Benjamini-Hochberg FDR correction across KS/Chi-Square tests, eliminating dashboard false alarms.
- Engineered an offline validation harness to stress-test detection thresholds against synthetic covariate and concept drift matrices, publishing verifiable true/false-positive metrics (<5% FP rate) to guarantee dashboard trust.

EDUCATION

B.Sc. Educational / Instructional Technology

Expected Graduation: 2029

Lagos State University | Lagos, Nigeria

CERTIFICATIONS & PROFESSIONAL DEVELOPMENT

- Machine Learning Specialization — DeepLearning.AI / Coursera (Andrew Ng)
- IBM Data Science Professional Certificate — IBM / Coursera
- TensorFlow Developer Certificate — Google
- Python for Data Science, AI & Development — IBM / Coursera
- Concepts of Artificial Intelligence & Machine Learning — Harvard Online
- Python for Data Science — Sololearn (Early Mastery Track)

ADDITIONAL INFORMATION

Availability: Immediate. Open to remote full-time, part-time, freelance contract, or internship roles globally.

Languages: English — Fluent (written and spoken).

Portfolio: solomonadenuga.cv | **GitHub:** github.com/TheLogeek | **LinkedIn:** [solomon-adenuga-6251a5316](https://www.linkedin.com/in/solomon-adenuga-6251a5316)