

Darlington Gospel

Senior Software & AI Engineer (Applied AI Systems)
Location: Port Harcourt, Nigeria | **Remote:** Available
Contact: darlontongospel@gmail.com | +234 8106056545
LinkedIn: [linkedin.com/in/darlington-gospel](https://www.linkedin.com/in/darlington-gospel) | **GitHub:** github.com/Daltonic
Website: dappmentors.org

Professional Summary

Senior Software & AI Engineer with 10+ years of experience and 6+ in production AI systems. Architected RAG pipelines, agentic workflows, and LLM apps across web/Discord/Telegram. Built end-to-end products from ingestion to response generation, integrating providers like OpenAI/Claude/Gemini/DeepSeek. Focused on scalable, cost-optimized services with real-time querying. Founded Dapp Mentors, educating 5,800+ via courses, 45+ articles, and 140+ videos.

Core Technical Skills

- AI & Applied Machine Learning**
- LLM Integration & Orchestration:** OpenAI, Claude (Anthropic), Gemini, DeepSeek, Ollama
 - Retrieval-Augmented Generation (RAG):** Vector databases (Chroma), semantic search, embedding pipelines
 - Agentic Systems:** LangChain, multi-agent workflows, autonomous decision-making
 - Data Processing:** Web scraping (Crawl4AI, BeautifulSoup), data ingestion pipelines, chunking strategies
 - AI System Design:** Prompt engineering, context management, cost optimization, latency reduction
 - Production AI:** Model deployment, monitoring, reliability engineering, A/B testing
- Backend & Infrastructure**
- Backend Frameworks:** Python (FastAPI, Flask), Node.js (Express)
 - Databases:** MongoDB, PostgreSQL, SQLite, vector databases
 - DevOps & Deployment:** Docker, GitHub Actions CI/CD, Linux/VPS management
 - API Design:** RESTful APIs, WebSocket connections, streaming responses
 - Async Processing:** Task queues, background jobs, real-time data pipelines
- Frontend & Interfaces**
- Web Frameworks:** React, Next.js, TypeScript
 - UI Development:** Tailwind CSS, Gradio interfaces, responsive design
 - Chat Platforms:** Discord bot development, Telegram bot APIs
 - AI-Driven UIs:** Streaming LLM outputs, conversational interfaces, real-time feedback

Professional Experience

Founder & Lead AI Engineer — Dapp Mentors / Daltonic LLC (US-based)

- 2022 - Present**
- Architected/deployed AI systems (RAG, multi-agents, LLM automation) for thousands
 - Built FastAPI backends with Chroma vectors, LangChain orchestration, multi-channel UIs
 - Designed ingestion pipelines for crawling/chunking/indexing web content
 - Implemented context-aware responses with prompt optimization and LLM routing
 - Owned lifecycle: design, Docker deploy, monitoring, metrics-based iteration
 - Created Dapp Mentors with AI features for 5,800+ learners
 - Published 45+ articles/140+ videos on LangChain/RAG/agents
 - Developed courses on AI bots, LangChain, ML applications

AI Data Evaluator (Contract) — AfterQuery AI Lab

- 2026 - Present**
- AI dataset curation and quality evaluation
 - Benchmarking LLM outputs against defined quality thresholds
 - Testing for data integrity and structural consistency
 - Crafting realistic, high-difficulty tasks designed to challenge frontier AI models
 - Evaluating task solvability and calibrated difficulty to ensure training data meets production-grade standards

Senior Software Engineer — Independent / Project-Based

- 2019 - 2022**
- Designed and implemented backend services supporting AI workflows, data-heavy applications, and real-time processing systems
 - Built production APIs and backend architectures using Python (FastAPI, Flask) and Node.js for ML-driven applications
 - Delivered systems with real-time data ingestion, asynchronous processing, and intelligent decision logic
 - Deployed and maintained applications using Docker, CI/CD pipelines, and VPS infrastructure

Technical Writer & Developer (Remote) — CometChat

- 2020 - 2022**
- Contributed to developer-facing integrations and tooling for real-time communication platforms
 - Built practical implementations alongside technical documentation for complex backend systems
 - Supported product teams by translating technical concepts into usable examples and integration guides

Software Developer (Full-Time) — GIT Technologies Ltd., Port Harcourt, Nigeria

- 2019 - 2020**
- Led development of full-stack software solutions with backend APIs and data processing systems
 - Built production systems using JavaScript, Python, and PHP with early AI/ML explorations
 - Trained teams and junior developers on backend development, frontend integration, and emerging AI fundamentals

Software Developer (Internship) — White Creativity, Port Harcourt, Nigeria

- 2017 - 2018**
- Developed web and mobile applications using React and Node.js with exposure to emerging AI concepts
 - Built scalable frontend and backend components with focus on performance and maintainability
 - Mentored junior developers on software development best practices and architectural patterns

Key Projects & Achievements

OpenBench
2026 — Built OpenBench, an AI dataset curation and evaluation platform with a 9-stage automated pipeline. Implemented CrewAI multi-agent probes (Task Analyst, Solution Architect, Code Implementer, Test Verifier) running in staggered parallel across 5 attempts per task, with early-exit logic to cut costs the moment an outcome is decided. Used LLM-as-judge scoring across 11 quality criteria and integrated OpenRouter as the sole LLM provider, routing across DeepSeek, Claude, GPT, Gemini, and Llama. Built Inngeist durable functions for crash-resistant pipeline execution across all 9 stages. Every approved training example and benchmark artifact is stored permanently on Sia Network via content-addressed storage.

Rebit
2026, in development — Building Rebit, a privacy-first, self-hostable video processing engine that uses local AI exclusively to transform long-form videos into short-form vertical clips. Whisper runs locally via FastAPI to produce timestamped transcripts from audio with no data leaving the user's infrastructure. An Ollama-powered LLM analyzes transcripts to detect high-value moments, hooks, and emotional peaks, ranking segments by virality score and hook strength. FFmpeg then maps selected segments back to exact video timestamps, cuts clips, converts them to 9:16 vertical format, and adds caption overlays. The entire pipeline runs through Inngeist durable functions with step-level status tracking. No proprietary cloud AI API is called at any point.

AI Content Production Engines
2026 — Built a personal suite of AI-powered content engines covering the full content production workflow: Content Forge for long-form blog generation, a strategic video script generator, a social post intelligence tool for LinkedIn, a thumbnail prompt engine, an architecture blueprint generator, a stack comparison reactor, a news asset forge, a 3D journey map generator, and an intelligence feed engine. Each tool integrates LLM APIs with structured prompt pipelines to produce consistent, high-quality output across formats. The suite now drives YouTube and LinkedIn content production at a pace and quality level that would be impossible to maintain manually.

RumbleCourt
2026 — Designed multi-agent debate system using OpenRouter LLMs (advocates/judge) for real-time reasoning. Built FastAPI backend for agent orchestration/streaming; integrated Opik SDK for tracing/evaluation. Added Solidity contracts via Hardhat for on-chain records. Created Next.js frontend with live interfaces. Featured in Commit To Change Hackathon for productivity gains.

AI Scraper Bot & Multi-Channel AI Assistant
2025 — Architected complete AI-powered knowledge retrieval system with web scraping and semantic search. Built FastAPI backend with automated crawling (Crawl4AI), document processing, and vector indexing (Chroma). Implemented full RAG pipeline: content ingestion, intelligent chunking, embedding generation, similarity search, and LLM-based response synthesis. Designed multi-channel deployment with Discord bot, Telegram bot, and Gradio web interface for real-time conversational access. Optimized for production with Docker deployment, GitHub Actions CI/CD, and Linux VPS hosting. System architecture enables community management teams to query knowledge bases instantly across multiple platforms with context-aware responses.

QuantumPool
2025 — Built AI-guided investment platform combining market data aggregation, sentiment analysis, and ML-driven decision workflows. Implemented FastAPI backend with real-time data processing, predictive modeling, and automated trading signals. Integrated with blockchain smart contracts for transparent fund management. Deployed complete production system with monitoring and performance optimization.

DappPay
2025 — Architected AI-driven payroll system with natural-language command processing. Built client-side LLM integration enabling users to execute payroll operations through conversational interfaces. Designed privacy-first architecture ensuring sensitive data never leaves client environment. Combined AI workflow automation with blockchain smart contracts for transparent, non-custodial payroll execution.

Project Descriptor
2025 — Developed Python-based project analysis tool using LangChain and OpenAI for intelligent codebase understanding. Implemented automated project decomposition breaking complex systems into analyzable components through advanced prompt engineering. Generates structured JSON outputs with architectural insights, dependency mappings, and improvement suggestions. Demonstrates practical application of prompt engineering and LLM orchestration for developer productivity tools.

MCP Servers (AI Note Management System)
2025 — Built AI-powered note management system with Gradio chat interface and SQLite backend. Implemented intelligent tagging, automated categorization, and smart search using LLM-driven workflows. Created conversational interface for natural-language note creation, retrieval, and organization. Demonstrates practical AI integration for personal productivity and knowledge management workflows.

Education

B.Sc. Computer Science
Ignatius Ajuru University of Education, Port Harcourt
2014 - 2018 | CGPA: 4.09 / 5.0

Open Source & Community

- Maintains multiple public repositories on GitHub focused on AI systems, RAG pipelines, LangChain implementations, and automation tools
- Regular contributor to applied AI and developer tooling projects with focus on production-ready implementations
- Leads Dapp Mentors community (5,800+ learners) providing free tutorials, open-source AI codebases, and educational content
- Published 45+ technical articles on AI crypto trading bots, automated AI systems, RAG architectures, and prompt engineering best practices
- Created 140+ YouTube videos covering LangChain development, AI agent workflows, vector databases, and production AI deployment

What Success Looks Like in My Work

- Production AI systems that are reliable, cost-efficient, and deliver consistent results at scale
- Clear architecture decisions balancing model capability, latency requirements, and operational costs
- Well-designed RAG pipelines where retrieval quality directly improves response accuracy
- Agentic systems that handle edge cases gracefully and degrade predictably when facing uncertainty
- Seamless integration across platforms (web, Discord, Telegram) with consistent user experience
- Comprehensive monitoring and logging enabling rapid debugging and continuous optimization