

ARON JOSE A

AI/ML · DevOps · Cloud Engineer

India · +91 95853 25511 · aronjose535@gmail.com · [LinkedIn](#) · [GitHub](#) · [Portfolio](#)

SUMMARY

B.Tech graduate (AI & Data Science) who builds and ships production-grade AI systems end to end — from LLM, RAG, and generative pipelines to the containerised, cloud-native infrastructure that runs them. Shipped CareerOS, a full-stack AI platform live on Kubernetes/AWS, and authored Flakehound, an open-source Python testing tool implementing spectrum-based fault localization (Ochiai algorithm) to pinpoint root causes of flaky tests. Strong across Python, FastAPI, LangChain/LangGraph, Docker, Kubernetes, Terraform, and CI/CD.

TECHNICAL SKILLS

Languages: Python, SQL, Bash/Shell Scripting

AI/ML & GenAI: TensorFlow, PyTorch, scikit-learn, GANs, CVAEs, CNNs, RNNs, LLM Fine-Tuning, RAG, Prompt Engineering, Agentic AI

LLM/Agent Frameworks: LangChain, LangGraph, LangSmith, CrewAI

Cloud: AWS (Cloud Architecting, CI/CD), Google Cloud Platform, Microsoft Azure (DP-900)

DevOps & Infrastructure: Docker, Kubernetes, Terraform, Jenkins, GitHub Actions, CI/CD, Git/GitHub, Linux, Apache Kafka, Redis, VPS, cPanel

Backend & Data: FastAPI, PostgreSQL, NumPy, Pandas, Matplotlib, Jupyter

Quantum ML: Qiskit, PennyLane (variational quantum circuits)

PROJECTS

CareerOS — Full-Stack AI Career Platform — [careeros.duckdns.org](#) ↗ **Live**

- Built an end-to-end AI platform orchestrating job search across role matching, tailored application generation, and pipeline tracking — React front end over a FastAPI service layer.
- Orchestrated multi-step LLM workflows with LangChain; PostgreSQL for durable state and Redis for caching and queues; integrated context-aware prompt engineering for personalised outputs.
- Containerised with Docker, deployed on Kubernetes, provisioned via Terraform on AWS with full CI/CD through GitHub Actions.

Tech: React, FastAPI, LangChain, PostgreSQL, Redis, Docker, Kubernetes, Terraform, AWS, GitHub Actions

Flakehound — Root-Cause Localization Tool for Flaky Python Tests — [github.com/aronjose2005/flakehound](#) ↗ **Open Source**

- Engineered a Python testing tool that identifies why a test is flaky — not just that it is — filling a gap where detection is commoditised but code-level diagnosis remains unsolved.
- Implemented spectrum-based fault localization (Ochiai algorithm), discriminating-variable analysis, static AST non-determinism scanning, and a two-phase tracing engine that avoids the Heisenbug observer effect (`sys.settrace` runs ~50x slower — so timing truth always comes from untraced phases).
- Ships as both a pytest plugin (`pytest --flakehound <nodeid>`) and a standalone CLI; zero runtime dependencies — pure Python standard library.

Tech: Python, sys.settrace, AST analysis, Ochiai SBFL, pytest plugin API

Hybrid Quantum-Classical GAN — Synthetic 12-Lead ECG Generation —

- Designed a quantum-classical GAN (QGAN) using variational quantum circuits to synthesise realistic 12-lead ECG waveforms from biomedical time-series data with GPU-accelerated adversarial training.
- Achieved stable adversarial convergence with realistic waveform generation; demonstrated quantum advantage potential for biomedical data synthesis.

Tech: PennyLane, PyTorch, FastAPI, NumPy, Python

EXPERIENCE

AI Research Intern · Tech Vaseegrah

May 2024 – Jul 2024

- Developed an AI-based e-commerce recommendation system in Python and scikit-learn, improving recommendation accuracy ~10% using collaborative filtering.
- Owned end-to-end data preprocessing, feature engineering, and model evaluation across the ML pipeline.

EDUCATION

B.Tech, Artificial Intelligence & Data Science

2022 – 2026

Karunya Institute of Technology and Sciences, Coimbatore

CERTIFICATIONS & LEADERSHIP

AWS: Academy Graduate – Cloud Architecting (2025); CI/CD Pipeline Fundamentals (2025) **Microsoft/NVIDIA:** Azure Data Fundamentals DP-900 (2024); NVIDIA AI for All – Basics to GenAI (2025)

Cisco/NPTel: Data Analytics Essentials; NPTel Cloud Computing (2024) **Forage:** Deloitte Data Analytics; TCS Data Analytics; GenAI-Powered Data Analytics (2025)

Leadership: President & Co-SPOC, Karunya-IIT Madras Incubation Cell Build Club (2024–25); Google Student Ambassador, Karunya Institute