

Marian Glen Louis

716-709-1834 | rockyglen86@gmail.com | www.linkedin.com/in/marian-glen-louis | glen-louis.vercel.app

SUMMARY

AI Engineer specialized in architecting autonomous, self-correcting agentic systems and multi-agent workflows using LangGraph and MCP. Expert in the full MLOps lifecycle, leveraging a background in Data Reliability to bridge the gap between legacy infrastructure and production-ready AI.

TECHNICAL SKILLS

- **Agentic AI:** LangGraph, Model Context Protocol (MCP), Agentic RAG, CRAG, Ragas (Evaluation), Groq.
- **MLOps & Cloud:** Terraform, AWS (IAM, S3, EC2), Docker, MLflow, DVC, GitHub Actions, Prefect.
- **Data Architecture:** Snowflake, Qdrant (Vector DB), Supabase, FastAPI, SQL, Python (PyArrow).
- **Machine Learning:** PyTorch, LightGBM, YOLOv8, OpenCV, Pandera (Data Contracts).

TECHNICAL PROJECTS

Aegis-Flow: Multi-Agent Cloud Security Orchestrator

- **Engineered an autonomous AI agentic workflow** using **LangGraph** to manage complex state transitions, including conditional routing for security auditing, automated remediation, and post-fix verification.
- **Developed a Model Context Protocol (MCP) server** to expose specialized AWS infrastructure tools (IAM, S3, EC2) to a Large Language Model, enabling precise, tool-augmented generation (TAG) and real-time environment interaction.
- **Architected a multi-agent system** featuring a specialized **Auditor** for risk detection and a **Remediator** for execution, separated by a **human-in-the-loop safety gate** to ensure alignment and operational security.
- **Implemented robust state management and persistence** using LangGraph's MemorySaver checkpointer, allowing the agent to maintain context across multi-turn security investigations and handle asynchronous human approvals.

AuditAI: Agentic RAG Compliance Engine

- **Engineered an Agentic RAG system** for automated compliance auditing, utilizing a multi-step reasoning engine to navigate complex regulatory documents and extract actionable insights.
- **Developed a high-performance ingestion pipeline** that processes unstructured data into a vector database, optimizing retrieval through advanced chunking strategies and metadata tagging.
- **Implemented a comprehensive LLM evaluation framework** using **Ragas**, tracking metrics such as faithfulness, answer relevance, and context precision to quantify and improve system performance.
- **Architected a scalable backend services layer** using FastAPI and Docker, ensuring robust API performance and seamless deployment for AI-driven compliance workflows.

F1-Apex Guardian | MLOps & AI Engineering

- **Automated CI/CD Pipeline:** Engineered a zero-touch MLOps workflow using **GitHub Actions** and **uv** to automate weekly telemetry ingestion, feature engineering, and model retraining for the 2026 season.
- **Cloud Feature Store:** Architected an **AWS S3** data layer and model registry, leveraging **joblib** and multi-layer caching to reduce dashboard latency by 85%.
- **Drift & Monitoring:** Integrated **Kolmogorov-Smirnov** statistical tests to detect feature drift in high-velocity telemetry, maintaining **Isolation Forest** accuracy via automated retraining triggers.
- **Deployment & Lineage:** Containerized a real-time **Streamlit** dashboard via **Docker** on Hugging Face, managing experiment lineage and versioning with **DagsHub** and **MLflow**.

Citi-Bike Taxi | Automated MLOps Pipeline

- **Architected an Automated MLOps Pipeline:** Developed a production-ready forecasting system using **Python** and **LightGBM** that bridges a 20-day historical data lag with a **Recursive Bridge** strategy to provide real-time hourly demand predictions.
- **Engineered Scalable Data Infrastructure:** Designed a disk-efficient streaming pipeline using **pyarrow** to process 12 months of NYC trip data (millions of rows) within GitHub Runner resource limits, utilizing **AWS S3** for feature and model storage.
- **Implemented Automated Model Governance:** Integrated **MLflow** for experiment tracking and **Evidently AI** for data drift monitoring, featuring an automated "Champion vs. Challenger" promotion logic to ensure model reliability and performance.
- **Developed Full-Stack Intelligence Dashboard:** Built a modern, responsive frontend using **Next.js** and **Tailwind CSS** that visualizes complex geospatial trip density and performance metrics through interactive **Recharts** and **Leaflet** maps.

WORK EXPERIENCE

Nissha Medical Technologies | Data Scientist Intern Capstone | Buffalo, NY.

Aug,2025- December 2025

- Developed an end-to-end vision system using **YOLOv8** and **OpenCV** to inspect 30M+ daily tickets, implementing automated defect detection.
- Transitioned the facility to a **predictive maintenance model**, reducing material waste and machine downtime via real-time defect alerts.

Wipro Technologies | Lead Data Reliability Engineer Bengaluru, India

May 2022 – Aug 2024

- **Automated Data Reliability Pipelines:** Engineered custom **Python-based validation frameworks** to automate the verification of large-scale data migrations from SQL Server to Snowflake, ensuring the integrity of datasets used for downstream analytics and machine learning.
- **ETL & Transformation Observability:** Collaborated with data modeling teams to programmatically verify complex business logic transformations, using **SQL and Python** to identify and resolve data drift and discrepancies during cloud-native transitions.
- **Technical Leadership & Code Review:** Directed a sub-team of 2 members, providing technical guidance in SQL development and requirement analysis while establishing standardized peer-review processes to ensure high-fidelity data deliverables.
- **Ground Truth Verification:** Established **Ground Truth benchmarks** by reconciling frontend Power BI metrics with backend Snowflake records, ensuring absolute data consistency—a critical requirement for preventing hallucinations in automated decision-making layers.

EDUCATION

Master of Science: Data Science, University at Buffalo, The State University of New York, December 2025

Bachelor of Technology: Electronics & Communication, VTU, India, June 2022

CERTIFICATIONS

- AI Engineering Core Track – Udemy
- AI Engineer Agentic Track: The Complete Agent & MCP Course – Udemy