Chowdry Mohammed Abdulla

AI/ML ENGINEER | AGENTIC AI | GCP | PRODUCTION AI SYSTEMS

+91 8618582225 | abdullachowdry@gmail.com | LinkedIn | Bangalore, India

PROFESSIONAL PROFILE

AI/ML Engineer and Data Science Lead with 5+ years of experience building production machine learning, Generative AI, agentic AI, and enterprise data engineering solutions. Experienced across the end-to-end lifecycle of AI systems—from data and retrieval pipelines to LLM applications, dynamic tool calling, API integration, deployment, monitoring, and production troubleshooting. Hands-on with Python, GCP, Vertex AI Gemini, LangGraph, FastAPI, Docker, Airflow, PostgreSQL, ClickHouse, BigQuery, Dataflow, and MLflow. Strong background in turning complex operational and data problems into reliable, scalable engineering solutions.

CAREER HIGHLIGHTS

- Built a production agentic AI chatbot for automated Apache Airflow ingestion-error investigation and root-cause analysis.
- Designed dynamic reasoning and tool-selection workflows integrating GCP/Dataflow, PostgreSQL, ClickHouse, semantic retrieval, and technical documentation.
- Built a multi-step agentic SQL optimization system using LangGraph, Ollama, BigQuery execution-plan context, and Vertex AI Gemini.
- Architected enterprise-scale Airflow ingestion and transformation platforms with data-quality validation across multi-tenant environments.
- Designed zero-downtime deployment strategies using active-copy and database-switching patterns.
- Re-architected derived-table execution with graph-based parent-child dependencies to enable parallel processing across hundreds of tables.
- Built ML deployment and serving pipelines using scikit-learn and MLflow for model versioning, reproducibility, and monitoring.

CORE TECHNICAL SKILLS

Agentic AI & GenAI	Generative AI, Agentic AI, LLM Applications, LangGraph, Ollama, Vertex AI Gemini, Function Calling, Multi-Step Reasoning, Prompt Engineering, Intent Classification, NER/Entity Extraction, Vector Embeddings, Semantic Search
AI/ML & MLOps	scikit-learn, ML Deployment & Serving, MLflow, Model Versioning, Model Monitoring, CI/CD for ML, Production AI Systems, Data Quality & Validation
Cloud & Data Engineering	Python, SQL, GCP, BigQuery, Google Dataflow, Apache Airflow, PostgreSQL, ClickHouse, Docker, FastAPI, REST APIs, Vector Databases
Engineering & Architecture	Enterprise Data Platforms, Multi-Tenant Systems, Workflow Orchestration, Root-Cause Analysis Automation, SQL Query Optimization, Zero-Downtime Deployment, Technical Documentation, Cross-functional Collaboration

PROFESSIONAL EXPERIENCE

Impact Big Data Analytics | Bangalore, India | Jan 2020 – Present

Progression: Data Science Lead (Jun 2024 – Present) | Senior Data Scientist (Sep 2023 – Jun 2024) | Data Scientist (Feb 2021 – Aug 2022) | Data Science Intern (Jan 2020 – Apr 2020)

Agentic AI Root-Cause Analysis — Airflow

- Built a production agentic AI chatbot that detects Apache Airflow ingestion errors and autonomously performs end-to-end root-cause analysis, replacing manual log triage.
- Implemented intent classification, entity extraction, vector-embedding semantic retrieval, and dynamic tool calling across errors, similar incidents, and technical documentation.
- Designed a dynamic reasoning loop where the LLM selects the next investigation tool based on context and intermediate findings rather than following a fixed workflow.
- Integrated investigation tools across GCP services, Dataflow jobs, PostgreSQL, and ClickHouse to analyze failures and generate structured remediation output.

Agentic SQL Query Optimizer Engine

- Designed and built a multi-step agentic architecture using LangGraph and Ollama to optimize SQL and support Text-to-SQL workflows.
- Implemented dynamic tool calling to identify the target SQL dialect and retrieve table schemas and, where available, BigQuery execution-plan context.
- Used retrieved schema and execution-plan context with Vertex AI Gemini to generate optimized SQL rewrites targeting improved query performance and cost.

Data Engineering & Platform Architecture

- Architected enterprise-scale Apache Airflow ingestion pipelines with automated standardization, validation, and data-quality checks across multi-tenant environments.
- Designed and maintained cloud-native data infrastructure using GCP, PostgreSQL, and ClickHouse for operational and analytical workloads.
- Designed multi-tenant database segregation supporting multiple databases and schemas per client.
- Extended analytical workloads to ClickHouse and supported data-platform architecture across OLTP and OLAP use cases.

Reliability, Scalability & Workflow Optimization

- Designed zero-downtime deployment using active-copy and database-switching strategies to reduce operational disruption during releases.
- Re-architected Derived Tables execution using a graph-based parent-child dependency framework, removing sequential processing dependencies and enabling parallel execution across hundreds of tables.
- Worked on production troubleshooting and root-cause analysis across ingestion, transformation, database, and ML workflows.
- Developed technical documentation and video walkthroughs used as onboarding resources for engineering teams.

ML Systems & Productionization

- Built automated ML deployment and serving pipelines using scikit-learn and MLflow for model versioning, reproducibility, and monitoring.
- Supported production ML workflows spanning model development, packaging, deployment, serving, and operational monitoring.

EDUCATION

Siddaganga Institute of Technology, Tumkur | 2016 – 2020

Bachelor of Engineering (BE) — Computer and Information Sciences | CGPA: 7.6/10