

# ADITYA SURESH

Sydney, NSW | P: +61 468 442 565 | hello@adityasuresh.com | [LinkedIn](#)

## PROFILE

Data Scientist with 3+ years of experience building scalable data platforms and production-grade analytics solutions across cloud environments, bridging data engineering and applied AI. Proficient in Python and SQL with experience developing data pipelines, model-driven workflows, and intelligent automation solutions on AWS and Azure platforms using modern MLOps and CI/CD practices. Recently specialized in Generative AI through a Master's in Artificial Intelligence, focusing on LLMs, Retrieval-Augmented Generation (RAG), and agentic architectures using frameworks such as LangChain and HuggingFace for real-world AI applications.

## WORK EXPERIENCE

### INFOSYS

Technopark, IN

Data Scientist

Nov 2020 - Aug 2023

- Selected among the top 1% of Infosys engineers (Power Programmer track); designed and delivered scalable analytics and AI-driven solutions using Python and SQL, building robust data processing and transformation workflows across Snowflake and AWS environments to support high-volume freight and operational datasets from multiple heterogeneous sources.
- Developed and deployed predictive and optimization models using Python (pandas, NumPy, scikit-learn) and R for routing efficiency, asset utilization, and anomaly detection; implemented model training and batch inference pipelines on AWS (S3, EC2, Lambda, SageMaker, Athena) to enable production-grade decision support systems.
- Engineered end-to-end ML workflows including feature engineering, model evaluation, hyperparameter tuning, and validation; utilized MLflow for experiment tracking and model versioning to ensure reproducibility and controlled deployment across environments.
- Operationalized ML systems using Docker and Kubernetes with CI/CD workflows (Jenkins, GitHub Actions) to automate testing, deployment, and monitoring; improved system reliability and reduced manual validation effort by 40% while maintaining SLA-driven reliability.

## EDUCATION

### UNIVERSITY OF NEW SOUTH WALES (UNSW)

Sydney, AU

Master of Information Technology (Artificial Intelligence), WAM: 81 (Excellence)

Sep 2023 - Aug 2025

### APJ ABDUL KALAM TECHNOLOGICAL UNIVERSITY

Trivandrum, IN

Bachelor of Technology in Computer Science & Engineering

Aug 2016 - Jul 2020

## PROJECTS

### Multi-Agentic Flight Booking Assistant | Python, FastAPI, LLMs (OpenAI), Agentic AI, MCP

Jan 2026 - Present

- Actively building and iterating an agent-driven conversational flight booking system using Python, FastAPI, and OpenAI LLMs, implementing structured intent detection, constraint validation, and guardrail-based requirement gating (origin, dates, budget, passengers, cabin class). Designed a chat-first hybrid UI (Next.js) with modular tool-calling architecture and Redis-backed session memory, reducing incomplete user queries by ~40% in simulated evaluation flows.
- Architecting a modular multi-agent orchestration framework (LangChain + MCP) supporting retrieval-augmented workflows, flight API integration, price monitoring, ranking heuristics, and alert pipelines. Incorporating vector-based retrieval, evaluation patterns (LLM-as-a-judge), and observability hooks to enable scalable, production-grade agent deployment in future iterations.

### Phenotype Concept Recognition Optimization | Python, Tensorflow, Scikit-learn, React, Flask

May 2025 - Aug 2025

- Designed and fine-tuned domain-specific NLP models (PhenoBERT, PhenoGPT) for extracting Human Phenotype Ontology (HPO) terms from unstructured biomedical text, building an end-to-end evaluation pipeline that improved F1-score by 10% and semantic similarity by 6% on a newly expanded gold-standard corpus.
- Developed scalable evaluation and visualization infrastructure using Flask and React, generating structured BioC XML outputs for seamless integration with annotation tools (INCEpTION) and supporting clinical genomics research workflows.

## SKILLS

**Programming & Data Science:** Python, SQL, pandas, NumPy, scikit-learn, TensorFlow, PyTorch, NLP, R

**Generative AI & Agentic Systems:** LLMs, RAG, AI Agents, LangChain, LangGraph, LlamaIndex, HuggingFace, OpenAI APIs, AWS Bedrock, Semantic Kernel, Pydantic AI, DSPy, MCP, AutoGen, Custom Embeddings, vLLM, LangSmith, TruLens

**Retrieval, Search & Vector Systems:** FAISS, Pinecone, ChromaDB, Weaviate, ElasticSearch, Hybrid Search (BM25 + Dense Retrieval)

**Cloud, MLOps & Deployment:** AWS (S3, Lambda, EC2, Athena, SageMaker), Vertex AI, Azure Synapse, Snowflake, MLflow, Weights & Biases, Kubeflow, Docker, Kubernetes, FastAPI, Redis, CI/CD Pipelines, RAGAS, Guardrails