

Ahmed Irfan

AI & Data Engineer | MLOps | Cloud & GenAI

ahmedbinirfan@gmail.com | Melbourne, VIC | www.linkedin.com/in/ahmedbinirfan/ | github.com/ahmedbinirfan

PROFESSIONAL SUMMARY

I'm an AI & Data Engineer with deep experience in MLOps, cloud infrastructure, and data analytics, turning complex models and pipelines into systems teams can rely on in production.

My toolkit spans Python, SQL, AWS, Kubeflow, and GenAI frameworks including LangChain and PyTorch. I've designed end-to-end MLOps pipelines, built RAG powered applications, and coordinated scalable enterprise data workflows.

I'm driven by autonomous AI agents and modern cloud platforms that automate intricate work and deliver clear, measurable outcomes for the business.

SKILLS

Languages & frameworks: Python, SQL, TypeScript, R, Next.js, Shell

AI & machine learning: PyTorch, scikit-learn, LangChain, LangGraph, AWS Bedrock, SageMaker

Cloud & Infrastructure: AWS EKS, EC2, S3, Lambda, Vertex AI, Athena

Data engineering: Databricks, Teradata, SQL Server, Pandas, Oracle, NumPy, MongoDB, AWS Glue, GitLab, Redshift, Spark, Kafka

Tools & platform: Power BI, Terraform, Vercel, QuickSight, Kubeflow, Supabase, Azure DevOps, Tableau, Jenkins, Vega-Lite, Git, Excel, KServe

PROFESSIONAL EXPERIENCE

Data Engineer

Feb 2026 – Present

Australian Taxation Office

Melbourne, VIC

End-to-end ML lifecycle on enterprise AWS.

- Managing the end-to-end machine learning lifecycle on an enterprise AWS platform, empowering 50+ data scientists to seamlessly transition from project onboarding and Kubeflow environment provisioning to full scale model training.
- Developing a Gen AI chatbot leveraging Retrieval-Augmented Generation (RAG) against comprehensive ATO knowledge bases, leveraging LangChain, AWS etc, automating responses to 10,000+ monthly queries and reducing manual query resolution time by 40%.
- Performing critical BAU operations and ML infrastructure maintenance, executing monthly Amazon EKS cluster patching and optimizing model inference pipelines via KServe and ModelServe to guarantee 99.9% platform uptime.
- Administrating centralized ML libraries and dependencies via Artifactory, ensuring secure, standardized package distribution for 100+ concurrent projects, which accelerated environment configuration times by 50%.

Forward Deployed Engineer (Weekend/Freelance)

Aug 2025 – Present

Webifex Labs

Data Analyst

May 2023 – Feb 2026

Australian Taxation Office

Melbourne, VIC

SQL analytics, Power BI, and model QA/UAT.

- Developed a complex SQL risk stratification model utilizing taxpayer financial data to segment income behaviours (non-lodgers/omitted income), assigning targeted risk tiers to drive strategic compliance interventions and uncover key trends.

- Architected interactive Power BI financial dashboards for 3 ATO business units, eliminating 10 hours of manual reporting weekly, while upskilling 2 junior data analysts to boost the team's overall reporting capacity by 25%.
- Spearheaded QA and User Acceptance Testing (UAT) for analytical financial models, collaborating with Data Science teams to execute rigorous test cases that identified critical anomalies and improved pre-deployment accuracy by an average of 15%.
- Optimized continuous financial data pipelines utilizing GitLab and Agile methodologies, bridging the gap between technical execution and diverse stakeholder requirements to deliver secure, scalable data solutions.

Data Engineer

Nov 2021 – May 2023

Civica Asia Pacific

Melbourne, VIC

Pipelines, SQL automation, and CI/CD.

- Orchestrated end-to-end data pipelines, encompassing data cleaning, transformation, and aggregation, and translated complex datasets into clear, actionable stakeholder dashboards utilizing Tableau and Excel.
- Architected and implemented 20+ SQL stored procedures to automate recurring data reconciliation tasks, reducing manual processing time by 40% and increasing data accuracy by 90%.
- Optimized CI/CD pipelines utilizing Azure DevOps, Git, and Jenkins, reducing system downtime by 25% and cutting feature deployment times by 30%, while enforcing strict coding standards through rigorous peer reviews.
- Mentored and upskilled 3 junior data engineers in SQL and best practices, cultivating a highly cohesive, technically proficient team.

Junior SQL Developer

Nov 2020 – Feb 2021

Civica Asia Pacific

Melbourne, VIC

Database apps and change documentation.

- Analyzed, researched and troubleshooted issues in database applications.
- CRUD operations and data manipulation skills with SQL. Design and implement solutions to avoid recurring problems.
- Used of Azure DevOps with GIT and Jenkins.
- Maintained evidence and documentation for all the changes proposed to recurring problems.

EDUCATION

Deakin University

Master of Applied Artificial Intelligence

Expected graduation July 2027.

Monash University

Bachelor of Computer Science in Data Science

CERTIFICATIONS

- Power BI: Developer — ATO Learning and Development
- AWS Certified Cloud Practitioner — AWS (Apr 2024)
- AWS Knowledge: Cloud Essentials — AWS (Feb 2024)
- Academy Accreditation - Databricks Lakehouse Fundamentals — Databricks (Jun 2023)
- AWS Cloud Concepts — DataCamp (Nov 2023)
- Introduction to DAX in Power BI — DataCamp (Oct 2023)

- Introduction to Power BI — DataCamp (Oct 2023)
- Building AI Browser Agents — DeepLearning.AI (Jun 2025)
- Introduction to Generative AI — Google Cloud (Feb 2024)
- Advanced RAG with Vector Databases and Retrievers — IBM
- Build RAG Applications: Get Started — IBM
- Develop Generative AI Applications: Get Started — IBM
- Vector Databases for RAG: An Introduction — IBM
- Docker Training Course for the Absolute Beginner — KodeKloud
- Kubernetes for the Absolute Beginners - Hands-on Tutorial — KodeKloud
- LangChain Essentials – Python — LangChain
- Quickstart: LangGraph Essentials - Python — LangChain
- Teradata - SQL — Teradata (Jun 2023)