Mohammad Mehdi Lotfinejad

Lead Data Scientist · Data & AI Engineer · HRDC-Accredited Corporate Trainer

+60 14 220 1751 · lotfinejad@gmail.com · linkedin.com/in/lotfinejad · github.com/datatweets · lotfinejad.ir

15+	1,000+	2TB+	20+
years in data & AI	professionals trained	telco data processed daily	production data pipelines

PROFESSIONAL SUMMARY

Data and AI leader with 15+ years of experience designing data platforms, building analytics, machine learning and LLM solutions, and delivering enterprise technical training. Currently leading agentic risk analytics at CIMB. Deep hands-on expertise in Apache Spark, Databricks, Airflow, AWS, RAG and agentic systems (LangChain, LangGraph), backed by DevOps, MLOps and LLMOps practice for production-grade reliability. Track record across banking, energy, telecom and legal tech.

KEY ACHIEVEMENTS

- Lead agentic risk analytics and model monitoring on Azure Fabric at **CIMB**, using LangChain, LangGraph and Rayfin to automate risk investigation and model performance oversight.
- Built and operated **20+ production pipelines** processing **2TB+ of telco data daily** at **99.9% availability** for analysts in 9 markets.
- Architected AI legal-data infrastructure for **10K+ daily case-analysis requests** and **50K+ documents per day**, cutting case turnaround from **48 to 18 hours**.
- Trained **1,000+ professionals** across ASEAN through HRDC-accredited programmes in data science, Python, cloud and big data.

CORE SKILLS

Data Science & AI	Machine Learning, NLP, LLM Applications, RAG, Agentic AI, LangChain, LangGraph, Rayfin, Prompt Engineering
Data Engineering	Apache Spark, PySpark, Airflow, Databricks, Apache Doris, ETL/ELT, Data Warehousing, Data Pipelines
Cloud Platforms	Azure, Microsoft Fabric, Azure OpenAI, AWS (Redshift, S3, Athena, Lambda, SageMaker), BigQuery, Snowflake
MLOps · LLMOps	Agent evaluation, model performance and drift monitoring, pipeline monitoring, CI/CD for data and AI
BI & Analytics	Power BI, DAX, Tableau, Dashboard Design, Data Modeling, Business Analytics
DevOps	Docker, Kubernetes, Helm, Terraform, Jenkins, GitHub Actions, Prometheus, Grafana, ELK
Programming & DBs	Python, SQL, PostgreSQL, MySQL, SQL Server, MongoDB, Neo4j, JavaScript/Node.js, Bash
Leadership	Corporate Training, Curriculum Design, Technical Writing, Team Leadership, Stakeholder Management

PROFESSIONAL EXPERIENCE

Lead Data Scientist

Aug 2026 – Present

CIMB · Kuala Lumpur, Malaysia

- Lead the design and delivery of agentic risk analytics and model monitoring solutions on Azure and Microsoft Fabric for credit, market and operational risk stakeholders across the group.
- Architect the Fabric data foundation for risk analytics: Lakehouse medallion layers, Data Factory pipelines and Power BI semantic models as the single source for risk reporting.
- Build multi-agent workflows with LangChain, LangGraph and Rayfin, orchestrating retrieval, tool calling and human-in-the-loop review to automate risk investigation, documentation and reporting.
- Deliver automated model performance monitoring covering stability (PSI/CSI), discrimination (Gini/KS), calibration and data drift, with threshold-based alerting to model owners.
- Establish LLMOps for agentic systems (prompt versioning, evaluation, guardrails, cost and latency monitoring, CI/CD) and align solutions with model risk management standards alongside validation and governance teams.

Principal Lead Data Scientist & Corporate Trainer

May 2023 – Aug 2026

The Center of Applied Data Science (CADS) · Kuala Lumpur, Malaysia

- Led data science, analytics and data engineering engagements for CIMB, Petronas, Shell and TNB, from proposal through deployment.
- Designed Spark-based ETL and analytics pipelines on AWS and GCP for enterprise reporting, data preparation and machine learning.
- Developed Power BI dashboards used by 40+ stakeholders across marketing, operations and finance for quarterly strategic decisions.
- Architected migrations from legacy on-premises systems to AWS, reducing infrastructure costs by about 30%; led a team of 5 data scientists and designed enterprise training programmes.

Chief Data & Knowledge Officer / AI Data Platform Lead

Feb 2024 – Feb 2025

Magna.ai (LawTech) · Florida, USA · Remote

- Owned the end-to-end AI legal data platform: data engineering, Neo4j graph analytics, Redshift warehousing, RAG-ready content processing and ingestion of 50K+ legal documents daily.
- Architected AWS infrastructure (Terraform, Helm) for 15+ microservices and Kubernetes services handling 10K+ daily case-analysis requests with auto-scaling.

- Automated legal case workflows, reducing manual intervention by 60% and cutting average turnaround from 48 to 18 hours.
- Implemented observability and governance: Prometheus, Grafana, ELK, least-privilege IAM, encryption and Vault.

Senior Data Engineer
Mar 2020 – Mar 2023

Axiata Digital Advertising (ADA) · Kuala Lumpur, Malaysia

Data analytics and AI subsidiary of Axiata Group, operating across 9 markets in South and Southeast Asia.

- Designed and operated 20+ production pipelines on Databricks and Spark, processing 2TB+ of telco data daily on multi-region AWS (Redshift, S3, Athena, Lambda) at 99.9% availability.
- Administered Snowflake and Redshift warehouses holding 20TB+ of data, improving query performance by 40% through table design and distribution-key tuning.
- Built Power BI and Tableau marketing analytics (Facebook Ads, Google Ads, proprietary sources) for 30+ analysts in 9 markets.
- Introduced Terraform provisioning, CI/CD and monitoring (Jenkins, CloudWatch, Grafana) for pipeline health and data quality; mentored 3 data engineers.

Senior Data Scientist & Corporate Trainer
Jun 2018 – Feb 2020

The Center of Applied Data Science (CADS) · Kuala Lumpur, Malaysia

- Bootstrapped the company's data engineering practice and designed multi-channel marketing analytics pipelines on AWS and GCP.
- Delivered data warehousing, ML and big data solutions for clients such as CIMB and TNB, and created and delivered HRDC-accredited programmes in data science, Python, cloud analytics and big data across ASEAN.

Big Data Engineering Team Lead
May 2016 – Jun 2018

Raha (Knowledge-Based Company) · Ahvaz, Iran

- Led 6 data engineers building Hadoop and PySpark pipelines processing 100GB+ daily on a 10-node cluster, feeding downstream recommendation systems.

EARLIER CAREER

Faculty Member & University Chancellor · Payame Noor University & Islamic Azad University, Iran — **2011 – 2018**

Senior Web Application Developer · Smarter Broadcast, California, USA · Remote — **2008 – 2011**

Business Analyst · Vira Co., Tehran, Iran — **2007 – 2008**

EDUCATION

PhD candidate, Computer Science · Universiti Malaya (UM), Malaysia — **2016 – Present**

MSc, Computer System Architecture · Islamic Azad University, Iran — **2007 – 2010**

BEng, Computer Hardware Engineering · Islamic Azad University, Iran — **2000 – 2004**

CERTIFICATIONS

Business CORe Program · Harvard Business School
Business Analytics, Economics for Managers, Financial Accounting

HRDC Accredited Trainer · HRD Corp, Malaysia

AWS Training & Certification · AWS
Technical Essentials, Practical Data Science with Amazon SageMaker, Big Data on AWS, Data Warehousing on AWS

PUBLICATIONS & COURSE AUTHORING

- Course author, *Window Functions in SQL* — Dataquest.
- Three books: *Object-Oriented Programming in C#*, *Project Management Fundamentals*, *Essentials of Data Mining*.
- 10+ peer-reviewed journal and conference papers in machine learning, neural networks, computer architecture and data mining.

LANGUAGES

English · Professional working proficiency

Persian · Native