

Ahmed Doghri

EXECUTIVE SUMMARY

Machine Learning Engineer with **7+ years** of experience developing advanced AI systems, specializing in NLP, LLMs, and predictive pipelines. Led high-impact projects in retail, pharma, and banking.

FUNCTIONAL EXPERTISE

- LLM Workflow Orchestration & Validation
- Deep Learning & NLP Research
- MLOps & Cloud Infrastructure Optimization
- Data Engineering & Real-Time Processing
- Cross-functional Communication
- Data Visualization

FUNCTIONAL EXPERTISE DEMONSTRATED

Weave | USA

03/2024 – 12/2024

Senior Data Scientist

Led the automation of regulatory submissions using LLMs.

- **LLM Workflow Orchestration & Validation:** Orchestrated over **7** complex LLM agent pipelines to fully automate the drafting of IND applications for FDA submissions. Implemented robust grounding techniques that decreased hallucination incidents by **60%**. [Python, LangChain, LlamaIndex, GPT-4o].
- **RAG System:** Built a Retrieval-Augmented Generation (RAG) system for precise table data extraction from regulatory documents. Integrated Elasticsearch and Spark, boosting search scalability and cutting parsing times by **40%**. [Elasticsearch, Spark, LangChain, PyTorch].
- **Model Optimization:** Streamlined inference workflows and applied model distillation to accelerate regulatory document classification, slashing prediction latencies by **25%**.
- **Cross-functional Collaboration:** Partnered with product managers and compliance teams to meet regulatory and security standards. Created documentation and knowledge transfer guides, boosting operations team productivity.

BenchSci | Canada

01/2022 – 03/2024

Senior Data Scientist (04/2023 - 03/2024)

Headed multiple initiatives to Advance Text Enrichment in the Biomedical Sector using Transformer models and LLMs.

- **LLM Implementation & Fine-tuning:** Led the implementation and fine-tuning of advanced LLMs such as GPT-4, PaLM2, and LLaMA2 to enhance risk detection and biomedical entity recognition. Achieved a **30%** increase in model accuracy and **25%** improvement in processing speed, optimizing outputs for biomedical applications. [Vertex AI, LangChain, MLFlow].
- **MLOps Implementation & Infrastructure Optimization:** Developed and deployed Python-based ML models using GCP's Vertex AI, integrating Vertex AI Pipelines and Kubeflow for efficient model orchestration and monitoring. Utilized MLFlow for experiment tracking and model versioning, optimizing cloud infrastructure for enhanced scalability and operational efficiency. [Kubernetes, Docker, Kubeflow].

Machine Learning Engineer (01/2022 - 04/2023)

Built machine learning models/systems that address a range of classification and text mining problems.

- **NLP Research:** Developed highly scalable transformer-based classifiers (F1 > **95%**) and tools leveraging machine learning, data regression, and rules-based models. [PyTorch, Vertex AI, Python].
- **Data Engineering & Real-Time Processing:** Developed and optimized real-time data streaming pipelines using Kafka and Flink on GCP, reducing data processing latency by **35%** and leveraging parallel environments such as distributed clusters and GPUs to improve system throughput by **20%**. [Kafka, Flink, GCP, Dataflow, Beam].

IBM | Canada

08/2019 – 01/2022

Senior Data Scientist & Consultant (05/2021 - 01/2022)

Drove go-to-market cloud initiatives and oversaw the projects' technical solutions and implementation.

- **NLP Research:** Prototyped and integrated an innovative chat logs annotation pipeline that reports on ~**100K** daily chats and masks any sensitive data in the logs. [Python, spaCy, PyTorch, Flask, Celery].
- **Client Management & Communication:** Led multiple client discussions to determine the business need, evaluate the technical approach, and report the solution to different executive stakeholders.

Data Scientist & Consultant (08/2019 - 05/2021)

Consulted on complex analytics and advanced machine learning methods to develop reusable assets for current and future business problems.

- **Deep Learning:** Implemented several unsupervised NLP deep learning techniques (KNN and SVD) to cluster **1.2M** chat logs and developed a chatbot solution using Watson Assistant (WA) in both English and French languages. [Python, TensorFlow, MongoDB, Redis].
- **Data Pipeline Engineering:** Tuned and optimized the modeling pipeline for a Canadian pension fund to achieve a **5%** better return vs. historical performance over the worst industry sectors' forecasts. [Spark, SQL, XGBoost, R].
- **Virtual Assistant Development:** Deployed three advanced chatbot solutions for the banking sector, driving \$4M in contact center automation savings. [PyTorch, Transformers, Watson Assistant].
- **Data Visualization:** Designed and implemented 8 dashboards for retail and banking clients, using platforms like Tableau, Adobe, IBM Cognos, and Shiny R to deliver actionable forecasting insights.

Ford Motor Company | Canada

05/2017 – 05/2019

Research Assistant in Embedded Software (Internship)

Generates hypotheses and proof of concept designs and performs experiments for validation.

- **Data Pipelining & Engineering:** Successfully carried the engine's misfire detection with **98%** precision using Wavelet Transform (WT), Principal Component Analysis (PCA), and RNN classifier. [MATLAB].
- **Linux Expertise:** Designed, built, and tested a Data Acquisition system (DAQ) that allows signal conditioning, sampling at **100MHz**, and processing of car sensory signals under the hood. [Linux, C].

EDUCATION

● **McMaster University | Canada**

05/2017 – 05/2019

Master of Science in Mechatronics – Summa Cum Laude
GPA: 12/12 (100%)

● **KAIST University | South Korea**

09/2012 – 02/2017

Bachelor in Electrical Engineering – Summa Cum Laude
GPA: 4.05/4.3 (97.22%)