

# MOEEN BAGHERI

AI Engineer – Toronto, Canada

[moinbag@gmail.com](mailto:moinbag@gmail.com) – [moeenbagheri.com](http://moeenbagheri.com) – [github.com/bagherig](https://github.com/bagherig) – [linkedin.com/in/moeenbagheri](https://linkedin.com/in/moeenbagheri)

## SUMMARY

I'm an AI Engineer with over 7 years of experience building AI-driven software systems. I hold a **Master of Science in Data Science and Analytics** from Ryerson University and specialize in enterprise-scale **Generative AI** and **LLM agentic systems** that streamline data access and decision-making.

As the Lead Engineer for BASF's *FOX Talk2Data* project, I designed and implemented an AI data assistant that generates and executes SQL queries from natural language inputs and analyzes the results, integrating multiple dynamic data sources, business rules, and knowledge bases. The agent is now used by over **3000** users across BASF and various operating divisions for analytics and reporting. The project earned **BASF's 2025 Digital Award** for its innovation, impact, and enterprise-wide adoption.

Beyond FOX, I have delivered end-to-end generative AI solutions such as *YT Copilot* and built production-grade computer vision and NLP systems at *Super GeoAI Technology* and *IBM*. My expertise spans **LLM agent design/development**, **prompt engineering**, **cloud deployment**, and **MLOps**, underpinned by a strong foundation in **data science** and **machine learning**. I'm passionate about leveraging AI to drive business impact and innovation.

## QUALIFICATIONS

**Enterprise AI Development:** Lead developer of **FOX Talk2Data**, BASF's enterprise AI assistant that converts natural language into governed SQL, integrating multi-system data and metadata for analytics and reporting.

**Generative AI Integration:** Extensive experience building **LangChain/LangGraph**-based agentic workflows on **Databricks** and **Azure OpenAI** for business intelligence and decision support.

**Leadership & Communication:** Led cross-functional engineering teams, mentored developers, and delivered live demos and executive presentations to align AI solutions with business priorities.

**Evaluation & Reliability Engineering:** Designed large-scale evaluation frameworks for SQL validation, regression testing, and runtime monitoring to ensure predictable LLM agent behavior across versions.

**DevOps & Cloud Systems:** Skilled with **Databricks**, **Azure**, **MLflow**, and **Docker**, including CI/CD automation and model deployment pipelines for production workloads.

## SKILLS

**Core Competencies:** Generative AI Agent Design & Development – Technical Leadership & Mentorship – Fullstack & API Development – Evaluation Frameworks & Testing – Stakeholder Communication

**Languages:** Python – SQL – JavaScript – Java – C# – R

**Tools & Technologies:** LangChain/LangGraph – Databricks – MLflow – Azure DevOps – Docker – Git – PostgreSQL – Unix

**Packages:** Pandas – NumPy – PySpark – PyTorch – TensorFlow – LightGBM – MLflow – Scikit-learn – Matplotlib – HuggingFace Transformers – OpenAI API – Databricks SDK – Azure SDK

**Concepts:** CI/CD – Agentic LLMs – RAG – Time-Series Forecasting – Finetuning – Transfer Learning – Computer Vision – Hyperparameter Optimization – NLP

## EDUCATION

### Master of Science: Data Science and Analytics

[Ryerson University](#) – Toronto, ON, Canada

2020

GPA: 4.33/4.33

- **Supervisor:** Professor Konstantinos Georgiou

- **Thesis:** Monthly Sales Forecasting for Walmart with Advanced Machine Learning Techniques

### Honours Bachelor of Science: Specialist: Bioinformatics; Major: Computer Science

[University of Toronto](#) – Toronto, ON, Canada

2017

- Graduated with **high distinction**; University of Toronto Scholar Award

- **Supervisor:** Professor Zhaolei Zhang

- **Thesis:** Detection of Genetic Variations in miRNA Primary Transcripts Using Statistical Analysis

# EXPERIENCE

<b>Senior AI Engineer</b> <i>BASF</i> – Toronto, ON, Canada	Oct 2024 – Present
<ul style="list-style-type: none"><li><b>LLM Agent for Data Analytics:</b> Led the end-to-end development of <i>FOX Talk2Data</i>, BASF's flagship AI data assistant used by over <b>3000</b> users. The agent was recognized with <b>BASF's 2025 Digital Award</b> for its innovation, impact, and enterprise-wide adoption.</li><li><b>Natural Language to SQL:</b> Designed an agent using <i>LangChain/LangGraph</i> that generates and executes <i>Databricks</i> SQL queries from a natural language question, and analyzes the results using RAG on market trends and management report documents, while enforcing custom business rules, and resolving ambiguities with tools.</li><li><b>Dynamic Data and Knowledge:</b> Implemented configurable data sources, custom business rules, and RAG knowledge bases, enabling easy and quick onboarding of new operating divisions.</li><li><b>Evaluation:</b> Built a custom evaluation framework with automated tests to track SQL correctness, latency, and regression-free releases, ensuring stable production behavior.</li></ul>	
<b>Generative AI Engineer</b> <i>YT Copilot</i> – Toronto, ON, Canada	Sep 2023 – Oct 2024
<ul style="list-style-type: none"><li><b>YouTube LLM Agent:</b> Developed a Chrome extension that skips promotional segments, summarizes videos, and enables semantic search over YouTube content using Generative AI.</li><li><b>Context Window Management:</b> Used a <i>MapReduce</i>-style chunking and aggregation pipeline to efficiently handle LLM context limitations for long YouTube videos.</li><li><b>Fullstack Development:</b> Built a JavaScript-based frontend integrated with YouTube's UI and a <i>Flask/PostgreSQL</i> backend for APIs, persistence, and subscription logic. Integrated <i>Ory Kratos</i> for user authentication and <i>Stripe</i> for subscription and payment handling.</li><li><b>Finetuning &amp; Optimization:</b> Improved response quality and reliability using tool calling, structured outputs, fine-tuning, and prompt engineering techniques such as chain of thought prompting.</li><li><b>Cloud Deployment &amp; Containerization:</b> Deployed the platform on <i>Azure Kubernetes Service (AKS)</i> with <i>Docker</i>, ensuring scalability and high availability.</li><li><b>CI/CD Pipeline:</b> Implemented a <i>CI/CD</i> pipeline with <i>Azure DevOps</i> for automated testing, deployment, and monitoring.</li></ul>	
<b>Machine Learning Engineer</b>  <i>Super GeoAI Technology Inc.</i> – Toronto, ON, Canada	Oct 2021 – Jul 2023
<ul style="list-style-type: none"><li><b>Computer Vision:</b> Managed a team of 7+ engineers to develop object detection models using <i>PyTorch</i> and <i>TensorFlow</i> for seed grading and crop yield estimation from phone and drone images, deployed via <i>Flask</i> applications and improving efficiency for agricultural partners by over <b>80%</b>.</li><li><b>Model Enhancement:</b> Applied transfer learning from datasets such as <i>COCO</i> and <i>Pascal VOC</i> and performed hyperparameter optimization with <i>Ray Tune</i>, improving accuracy by <b>10%</b> for seed grading and <b>15%</b> for wheat head detection models.</li><li><b>System Design:</b> Built a <i>NextCloud</i>-based storage platform integrated with <i>Ory Kratos</i> authentication on Linux infrastructure, providing secure and reliable data and model access.</li><li><b>Leadership:</b> Conducted technical interviews, led team meetings, and enforced high-quality implementation standards through structured code reviews.</li></ul>	
<b>Machine Learning Engineer</b>  <i>IBM</i> – Toronto, ON, Canada	Jul 2021 – Jan 2022
<ul style="list-style-type: none"><li><b>LLM Design:</b> Used <i>BERT</i>-based models to perform NLP on Software Requirements Specification (SRS) documents to detect conflicting and duplicate requirements.</li><li><b>Model Optimization:</b> Achieved a <b>20%</b> increase in F1-score through fine-tuning and hyperparameter optimization of a Microsoft pre-trained model.</li></ul>	

## Instructor – Practical Deep Learning (CIND850) ↗

Sep 2021 – Dec 2021

Ryerson University – Toronto, ON, Canada

- **Mentorship:** Mentored 20+ graduate students on applied deep learning projects, reviewing model architectures and experimental design.
- **Assessment:** Graded assignments and exams and provided targeted feedback on deep learning implementations and results.
- **Curriculum Development:** Developed lab exercises and solutions emphasizing end-to-end deep learning workflows from data preparation to evaluation.
- **Lecture Delivery:** Delivered lectures and tutorials on core deep learning concepts, ensuring clarity and practical relevance.

## Quantitative Data Scientist

Jun 2014 – Aug 2019

Electronica AI – Toronto, ON, Canada

- **Trading Algorithm Development:** Collaborated with 5+ engineers to design and implement high-frequency trading algorithms with a daily order flow of **\$10M+**, leveraging machine learning to improve predictive accuracy.

## Quantitative Data Analyst

Aug 2016 – Oct 2018

Alkemi Network – Toronto, ON, Canada

- **Validation and Stress Testing:** Validated automated financial trading strategies based on intra- and inter-exchange price discrepancies and conducted stress testing on exchanges to identify bottlenecks and vulnerabilities in live-like scenarios.

## PROJECTS

### Wikipedia Article Question-Answering System

Aug 2024 – Sep 2024

Personal Project – Remote

- **Natural Language Processing:** Developed an NLP pipeline over Wikipedia articles for Named Entity Recognition (NER) and Question Answering (QA) using *LangChain*, *HuggingFace Transformers*, and *spaCy*.
- **Similarity Search & Embedding Generation:** Implemented *FAISS*-based similarity search and embeddings to retrieve relevant text chunks efficiently from large corpora.
- **System Deployment:** Deployed the application using *Flask* and *Docker*, with unit tests to ensure reliability of the QA pipeline.

### Natural Language SQL Querying System

Jul 2024 – Aug 2024

Personal Project – Remote

- **NLP to SQL Translation:** Built a natural language interface that converts user queries into SQL commands using *LangChain* and *HuggingFace Transformers*, allowing complex database operations without direct SQL knowledge.
- **Database Management & Integration:** Used *SQLAlchemy* with *PostgreSQL* for schema management, migrations, and query execution.
- **API Development:** Implemented RESTful APIs in *Flask* for natural language querying, with integrated tests for correctness and performance under varied workloads.

### Alzheimer's Disease Detection

Jul 2021 – Oct 2021

Ryerson University – Toronto, ON, Canada

- **Deep Learning:** Developed a deep learning model for Alzheimer's disease detection from MRI scans using *TensorFlow* and *Keras*.
- **Image Classification:** Used architectures such as *DenseNet121*, *ResNet50*, and *VGG16* to develop a deep learning model for Alzheimer's disease detection from MRI scans.
- **Data Augmentation:** Leveraged datasets including *OASIS* and *ADNI*, with preprocessing and augmentation to improve generalization.
- **Model Optimization:** Applied transfer learning and Bayesian hyperparameter optimization to achieve a notable **10%** accuracy boost in Alzheimer's disease detection.

## Coleridge Initiative

Mar 2021 – Jun 2021

*Kaggle Competition* – Remote

- **Natural Language Processing:** Developed a *BERT*-based text-to-text transformer model in *PyTorch* with *CRF* for extracting data citations.
- **Data Preprocessing:** Used NLP techniques with *spaCy* and *Gensim* for tokenization, stemming, and stopword removal.
- **Competition Achievement:** Achieved a top **5%** rank, earning a *silver medal* in the competition.

## Riiid! Answer Correctness Prediction

Oct 2020 – Jan 2021

*Kaggle Competition* – Remote

- **Time-Series:** Developed transformer and *LightGBM* models to predict whether a student will answer a question correctly based on their historical performance.
- **Model Efficacy:** Attained an AUC score of **80%**, demonstrating strong predictive performance.
- **Competition Achievement:** Achieved a top **3%** position in the Kaggle competition, earning a *silver medal*.

## Walmart Sales Forecasting

May 2020 – Aug 2020

*Ryerson University* – Toronto, ON, Canada

- **Time-Series:** Developed *LSTM* and *LightGBM* models to forecast 28-day sales for 3,049 Walmart products using historical and explanatory features.
- **Feature Engineering:** Engineered additional features to capture seasonal patterns and store-level effects, improving forecast quality.
- **Data Analysis:** Analyzed M5 competition data spanning ten Walmart stores across CA, TX, and WI from 2011 to 2016.
- **Optimization:** Applied *Bayesian optimization* to fine-tune model hyperparameters.

## Traffic Sign Recognition

Mar 2020 – Apr 2020

*Ryerson University* – Toronto, ON, Canada

- **Image Classification:** Designed and trained a *Convolutional Neural Network (CNN)* to classify German traffic signs, achieving **98.14%** accuracy on Google Maps Street View images.
- **Data Augmentation:** Implemented data augmentation techniques such as zoom, shear, and rotation to improve robustness and generalization.
- **Comparative Analysis:** Compared CNN performance against *ANN*, *Random Forest (RF)*, and *Support Vector Machine (SVM)* baselines.
- **EDA & Model Evaluation:** Performed exploratory data analysis and evaluated models using precision-recall curves, heatmaps, and loss curves.

## Genetic variation of miRNA primary transcripts in different ethnicities

Apr 2017 – Aug 2017

*University of Toronto* – Toronto, ON, Canada

- **Genetic Study:** Investigated genetic differences in miRNA primary transcripts across human populations, analyzing 1,127 SNPs from 1,881 miRNA transcripts to identify key trends.
- **Data Analysis:** Used variance and *Fixation Index (F-ST)* to quantify genetic differentiation, comparing “average of ratios” and “ratio of averages” formulations for genome-wide F-ST.
- **Genetic Research:** Showed that African populations exhibit the highest degree of genetic variation, supported by dendrogram-based analyses.