

Darshan Rahul Rajopadhye

(857)-376-9898 | rajopadhye.d@northeastern.edu | [linkedin.com/in/darshanrr](https://www.linkedin.com/in/darshanrr) | github.com/therrshan |

Available Full-time: June 2025.

Education

Northeastern University	Boston, MA
Khoury College of Computer Sciences	May 2025
<i>Master of Science in Artificial Intelligence</i>	3.96 GPA
<i>Relevant Coursework : Machine Learning, Information Retrieval, Data Mining, Algorithms</i>	
<i>Teaching Assistant : Machine Learning II (Prof. Chieh Wu)</i>	

Technical Skills

Languages & Frameworks	Python, Bash, Java, C++, SQL, FastAPI, Flask, Django, Pytest, Poetry, Streamlit, Pydantic, Jupyter
Infrastructure & DevOps	Kubernetes, OpenShift, Docker, GitHub Actions, Helm, Argo CD, GCP, AWS, CI/CD
AI & ML Tools	TensorFlow, PyTorch, Scikit-learn, MLflow, Airflow, DVC, LangChain, KServe, Hugging Face
Data & Databases	MySQL, PostgreSQL, MongoDB, BigQuery, Pinecone, Neo4j, Vector DBs (FAISS, Qdrant)
Streaming & Real-Time	Apache Kafka, Spark Streaming, Apache Flink, Celery, Pub/Sub, Redis Streams

Experience

Data Engineer	Aug 2021 – Sept 2022
Larsen and Toubro Infotech	Pune, IN
<ul style="list-style-type: none">Designed and optimized scalable ETL pipelines to process over 70GB of financial data daily, integrating structured and unstructured data to support advanced analytics and machine learning workflows.Developed automated data preprocessing and transformation pipelines using Python and Bash scripts, reducing manual intervention and enabling efficient feature extraction for downstream decision making.Designed high-performance storage solutions using AVRO and Parquet, enhancing data retrieval speeds.	

Projects

Credit Risk Explainer with Real-Time Fairness & Interpretability <i>Python, SHAP, LIME, oost, Link</i>	May 2025 – Jun 2025
<ul style="list-style-type: none">Built an interactive, explainable AI dashboard for real-time credit risk prediction using XGBoost, Random Forest, and Logistic Regression, with full model evaluation and comparison.Implemented local interpretability with SHAP and LIME to provide per-user prediction insights, promoting transparency and trust in AI-driven decision systems.Designed fairness monitoring modules with disparate impact and equal opportunity metrics, supporting responsible AI and bias detection across user segments.Deployed a Streamlit dashboard on Streamlit Cloud with real-time user input processing, automated model retraining, and CI/CD deployment via GitHub Actions for continuous delivery.	
E-commerce Churn Prediction with Real-Time Inference <i>Python, FastAPI, MLflow, Docker, GitOps</i>	Jan 2025 – Mar 2025
<ul style="list-style-type: none">Engineered a real-time ML pipeline with MLflow for experiment tracking, version control, and continuous integration, inspired by GitOps practices for production-ready deployment.Optimized XGBoost models via Bayesian search, achieving 92.85% accuracy and 0.94 AUC-ROC, while monitoring model fairness, bias drift, and demographic parity.Built a low-latency inference service using FastAPI with Pydantic-based schema validation, integrated content safety filters, and deployed as a containerized microservice on OpenShift.Architected a production-grade system using Docker Compose, with multi-container orchestration, automated health checks, and robust CI/CD triggers for real-time deployment.	
Predictive Analytics for NBA Player Performance and Game Outcomes <i>Python, scikit-learn, NBA-API</i>	Nov 2024 – Dec 2024
<ul style="list-style-type: none">Designed an end-to-end machine learning pipeline for predicting NBA game outcomes using real-time and historical player behavior data, achieving 83% accuracy.Engineered 30+ advanced features from raw NBA API streams, modeling individual player dynamics, fatigue, and team synergy for outcome forecasting.Evaluated and tuned models including XGBoost, LSTM neural networks, and ensemble methods to maximize AUC-ROC, with a 12% performance lift.Applied time-series analysis to monitor player fatigue and injury risk signals, enhancing model personalization and temporal prediction accuracy by 7%.	