

# RAHUL SHARMA

College Park, Maryland | +1-240-610-7281 | [rahul11s@umd.edu](mailto:rahul11s@umd.edu) | [@LinkedIn](#) | [@Git](#) | [@Portfolio](#)

## Summary

Data Scientist with 5+ years of experience applying statistical modeling & ML/DL to large-scale problems. Proven track record across NLP, CV, recommender systems, and time-series modeling driving impact in healthcare, risk, and operational analytics.

## Education

### University of Maryland - College Park, US

May 2026

*Master of Science in Data Science*

- Relevant Courses: Probability & Statistics, Machine Learning, Data Science, Algorithms, Data Modeling, Big-Data Systems, Computer Vision, Time-Series, Cloud Computing

### Indian Institute of Technology (BHU) Varanasi, India

July 2020

*Bachelor of Technology in Engineering*

- Relevant Courses: Eng. Mathematics, Computer Programming, Numerical Methods, Applied Mathematics.

## Work Experiences

### Bilbo.ai, Baltimore, US

Sept 2025-Present

*AI Engineer Intern*

- Designed and deployed scalable **RAG** system for unstructured medical documents; built automated data ingestion, embedding, evaluation, and inference pipelines with production-ready serving and performance monitoring.

### Jet2 Travel Technology, Pune, India

June 2022- Aug 2024

*Data Scientist*

- Designed and automated distributed ML training and inference pipelines for 1M+ transcripts using supervised fine-tuning and inference strategies; deployed scalable solutions enabling risk and operational decision support (**88% Jaccard recall**).
- Developed clustering, anomaly detection, and statistical evaluation frameworks with defined validation thresholds and performance diagnostics, improving detection specificity to **92%**.
- Operationalized LSTM/GRU forecasting models with automated retraining, walk-forward validation, **drift monitoring (<2%)**, and dashboard reporting to ensure reliable model lifecycle management.

### Fibe (Formerly EarlySalary), Pune, India

Sept 2021-June 2022

*Data Scientist*

- Led end-to-end credit risk and underwriting model development (feature engineering→training→validation→ deployment), improving ROC-AUC to **0.94** and integrating outputs into production systems, influencing credit and marketing strategies.
- Automated reproducible ML pipelines with added governance checks (data integrity, security), shrinking turnaround **15x**.

### Advanced Technology Consulting Service, Jaipur, India

Nov 2020-Sept 2021

*Associate Data Scientist*

- Engineered PySpark to ETL pipelines to structure large-scale production datasets, enabling downstream ML workflows.
- Developed and deployed computer vision inference (Mask R-CNN) on **Azure-backed** infrastructure, packaging models for scalable production serving at **85% mAP**.

### Amazon Pay India, Bangalore, India

May 2020-Nov 2020

*Research Intern*

- Built fraud-risk prediction model (**95% F1**) using clustering and statistical evaluation; informed fraud prevention initiatives.
- Designed OCF executive dashboards, translating analytical findings into actionable credit and marketing insights.

## Projects

- **MIMIC-CXR Radiology Report Generation:** Fine-tuned **multimodal foundation model (LLaVA-NeXT)** with **LoRA** to generate structured JSON reports from images and metadata inputs. Designed performance comparison and conducted **A/B evaluation** conditioning to validate model robustness and reasoning quality.
- **Multimodal Fashion Recommendation System, Fall'24, UMCP:** Built large-scale multimodal recommendation system leveraging Fine-tuned CLIP encoders and FAISS retrieval, improving text-image alignment and ranking performance (88% MOS) through iterative model evaluation and tuning.

## Publications

- **Co-author, IEEE ICSCDS 2025.** "Deep Learning for Reliable Routing in WANETs." *Peer-reviewed paper.* [Link](#)

## Skills

- **Programming:** Python, SQL, PySpark, C/C++,
- **ML & Modeling:** Statistical Modeling, Clustering, Anomaly Detection, Time-Series, Recommendation Systems.
- **Model Engineering:** ML Lifecycle, Model Validation, Performance Tuning, Drift Monitoring, Automated Retraining
- **DL & Generative AI:** PyTorch, TensorFlow, Transformers, LLMs, RAG, Fine-Tuning, Multimodal Models
- **Cloud, Big Data & MLOps:** AWS, Spark, GCP, Azure, Snowflake, Databricks, Docker, Tableau, Power BI, Dataiku