

HARSH SONI

United States | harshsoni@workmailit.com | (210) 201-5688 | [LinkedIn](#) | [GitHub](#) | [Portfolio](#)

SUMMARY

Data Scientist with 5+ years of experience architecting and deploying scalable machine learning systems across healthcare and enterprise environments, processing multi-terabyte datasets and delivering measurable business impact. Specialized in predictive modeling, NLP with transformer architectures, and end-to-end MLOps lifecycle, driving risk optimization, fraud detection, and data-driven decision systems in regulated domains.

SKILLS

Programming: Python, R, SQL

Frameworks & Libraries: Scikit-learn, TensorFlow, PyTorch, XGBoost, LightGBM, Pandas, NumPy, PySpark, Dask

Statistical & Analytical Techniques: Hypothesis Testing, Time Series Forecasting, Bayesian Inference

Machine Learning & Modeling: Supervised Learning, Unsupervised Learning, Feature Engineering, Hyperparameter Optimization

Deep Learning & NLP: Neural Networks, CNN, LSTM, Transformers, Large Language Models, Named Entity Recognition

Databases: PostgreSQL, MySQL, MongoDB, Snowflake

Data Visualization: Tableau, Power BI, Plotly, Matplotlib, Seaborn

Big Data: Apache Spark, Hadoop, Hive

MLOps & Deployment: MLflow, Docker, Kubernetes, CI/CD Pipelines, Model Monitoring, REST APIs

Cloud Platforms: AWS (S3, EC2, SageMaker), Azure ML, Google Cloud Platform (Vertex AI, BigQuery)

PROFESSIONAL EXPERIENCE

AIG, United States

January 2024 – Present

Data Scientist, AI

- Led development of gradient boosting risk adjustment models on claims and clinical datasets, increasing scoring precision by 28% and strengthening underwriting strategies for value-based care programs.
- Orchestrated PySpark and Snowflake data pipelines processing over 2TB of healthcare data daily, minimizing feature engineering latency by 37% and improving model training efficiency across analytics workflows.
- Designed transformer-based NLP solutions to extract structured medical entities from unstructured clinical notes, boosting extraction accuracy by 26% and supporting advanced clinical decision support applications.
- Established MLflow experimentation workflows with CI/CD integration, shortening model deployment timelines by 43% while ensuring regulatory compliance, audit traceability, and lifecycle governance across production systems.
- Built unsupervised anomaly detection models for fraud analytics, increasing detection precision by 24% and reducing financial exposure by identifying irregular billing behaviors within large-scale insurance claim systems.
- Implemented containerized inference endpoints using Docker and Kubernetes, enabling scalable batch and real-time predictions while maintaining system reliability, fault tolerance, and consistent performance across cloud environments.

Adons Softech, India

January 2019 – July 2022

Machine Learning Engineer

- Developed recommendation systems leveraging collaborative filtering and ranking algorithms, increasing customer engagement by 25% and improving personalization across high-traffic digital platforms serving diverse user segments.
- Constructed distributed data processing pipelines using Apache Spark and Hadoop ecosystems, handling datasets exceeding 600GB and enhancing transformation throughput by 39% for large-scale analytics operations.
- Created classification models for churn prediction using ensemble techniques, improving targeting accuracy by 27% and enabling business teams to implement proactive customer retention and engagement strategies.
- Engineered time series forecasting pipelines for revenue prediction, increasing forecast accuracy by 33% and supporting data-driven planning decisions across finance, operations, and executive leadership teams.
- Automated data ingestion, validation, and preprocessing pipelines using Python-based frameworks, improving data reliability by 44% and minimizing inconsistencies impacting downstream machine learning model performance.
- Delivered interactive business intelligence dashboards using Power BI and Plotly, improving stakeholder visibility into key performance indicators and accelerating reporting cycles across multiple operational business units.

EDUCATION

Master of Science in Data Science

August 2022 – May 2024

The University of Texas at Arlington, Arlington, Texas

Bachelor of Technology in Information Technology

June 2015 – May 2019

Gujarat Technological University, India

CERTIFICATIONS

- [DP-100: Designing and Implementing a Data Science Solution on Azure](#)
- [ICSI | CNSS Certified Network Security Specialist](#)
- [Completed 5-Day Gen AI Intensive Badge](#)