Abhir Rajesh Iyer

Chicago, IL | +1 (930) 904-4269 | abhir.iyer@gmail.com | abhir.com | linkedin.com/in/abhir-iyer

Summary

Data scientist (MS in Data Science, 2026) who applies causal inference, statistical modeling, and LLM evaluation to large datasets; now building data-quality metrics and investigating data anomalies in Snowflake at Northern Trust.

Skills

Statistics & Causal Inference: hypothesis testing, difference-in-differences and panel models (two-way fixed effects), logistic regression (MLE), robust standard errors, cross-validation, inter-rater agreement (Krippendorff's alpha), anomaly detection.

Machine Learning & AI: XGBoost, SHAP, PyTorch, Hugging Face, topic modeling (BERTopic), LLM evaluation (LLM-as-a-Judge, human agreement), LLM agents (OpenAI Agents SDK, LangGraph), AI-assisted development (Claude).

Data & Analytics: SQL (Snowflake, PostgreSQL, BigQuery), Python (pandas, NumPy), R, Spark, Databricks (MLflow, Unity Catalog), data-quality metrics, data visualization.

Engineering: FastAPI, Docker, Git, GitHub Actions CI/CD, AWS (S3, SageMaker), Linux, SLURM / HPC.

Certifications: SnowPro Core (Snowflake); Sigma Delivery Fundamentals, Certified Partner (Sigma Computing).

Experience

Hexaware Technologies September 2026 - Present

Data & AI Associate Engineer (Client: Northern Trust) *Chicago, IL*

- Built automated data-quality metrics and alerts in SQL for Northern Trust's Snowflake data hub, measuring delivery timeliness, duplicate keys, and weekend completeness across 5 investment position feeds.
- While building the weekend-completeness metric, found a silent gap that had dropped ∼34K production position rows over 30 holiday weekends since January 2025; traced the root cause and built a duplicate-free fix.
- Traced 25–53 minute nightly table refreshes to warehouse contention (8–16 concurrent queries) rather than query logic, and found that batch-triggered refreshes never appear in refresh history, a monitoring blind spot.

Data Engineering Intern (Databricks Center of Excellence) *June 2026 - August 2026*

- Built a multi-agent data-profiling tool on Databricks (OpenAI Agents SDK): agents profile tables, propose data models, and turn measured statistics into data-quality rules, with deterministic SQL checks to limit LLM cost and errors.
- Built docsig, an AI service that verifies signatures and dates on financial documents through a 4-tier cascade: digital signatures, text layer, OCR/CV, then a vision LLM on cropped regions only.
- Evaluated docsig on real documents and diagnosed failure modes such as OCR noise read as signer names, phantom signatures, and expired documents marked complete, fixing each with targeted rules and tests (92 in total).

Indiana University Bloomington January 2026 - May 2026

Graduate Research Assistant (FADS Fellowship) *Bloomington, IN*

- Built a multimodal LLM pipeline (WhisperX, Qwen2.5-32B on vLLM, Qwen2.5-VL) that codes TV episodes on 42 variables for a 50-year longitudinal media study: 1M+ records on an HPC cluster, replacing ∼15,000 hours of manual expert coding.
- Evaluated LLM output against human expert coders with Krippendorff's alpha, and built fault-tolerant batch resumption so multi-terabyte jobs restart without redoing finished work.

Projects

byltris.com - Consumer Financial Intelligence Platform

- Built a difference-in-differences panel model (two-way fixed effects) on 677,037 CFPB complaints that estimated an 11.8% post-merger rise in consumer complaints at Truist ($p = 0.014$).
- Found 14 complaint themes with BERTopic topic modeling, including a 302-complaint SunTrust migration cluster mostly from 2022–2023, and published all findings as an interactive site at byltris.com.
- Trained a bank-risk model on 550,404 FDIC financial records that flagged at-risk institutions out of sample (0.80 AUC, 100x precision lift), using model explainability to isolate the dominant risk driver.
- Estimated the adjusted gap in mortgage-approval odds across 4.49M HMDA applications with logistic regression.

TerraStream - End-to-End Computer Vision System for Automated Building Detection

- Fine-tuned a DINOv2 vision transformer with LoRA, training only 0.57% of its 87M parameters, to reach 0.85 mIoU on building segmentation, with a multimodal LLM-as-a-Judge (Llama 4 Scout) grading uncertain predictions.

Education

Indiana University Bloomington | **MS in Data Science** August 2024 - May 2026

Coursework: Social Media Mining, Applied ML, Advanced Databases *GPA: 3.60 / 4.0*

University of Mumbai | **B.Tech in AI & Data Science** December 2020 - May 2024

Coursework: Deep Learning, Neural Networks, Computer Vision *GPA: 3.67 / 4.0*