# Aman Singh

862-331-1884 | amanksingh0611@gmail.com | New York, NY | LinkedIn | Portfolio | GitHub

## WORK EXPERIENCE

**Community Dreams Foundation** August 2025 – August 2026
**Data Scientist** Remote, United States

- Quantified demand sensitivity across product categories with a **price-elasticity model** in Python, identifying pricing changes projected to increase net fundraising revenue by 8%.
- Lowered LLM inference cost 60% by routing over half of requests to lightweight models, after a **natural experiment** on rate-limit fallbacks showed nearly equal task completion.
- Engineered an evaluation layer that validates assistant answers against a 200+ pair golden dataset in CI, with LLM-as-a-Judge scoring and human escalation, reaching ~98% answer accuracy.

**NSquare Xperts** January 2022 – August 2023
**Data Scientist** India

- Analyzed application and transaction patterns for a consumer-financing client to flag elevated-risk segments and emerging fraud anomalies, turning them into **monitoring rules**.
- Reduced manual ticket triage 40% by fine-tuning DistilBERT for natural language processing (NLP) intent classification on customer complaint text, deploying it to production on **Amazon SageMaker**.
- Estimated heterogeneous treatment effects with a gradient-boosted **XGBoost T-learner**, guiding $300K+ in marketing spend toward customer segments with the highest predicted incremental conversion lift.
- Prioritized prospects with a conversion **propensity model** built on CRM and behavioral data, helping sales focus outreach across $2M+ in qualified pipeline.

**New Jersey Institute of Technology** September 2024 – May 2025
**Graduate Project Analyst** Newark, NJ

- Matched 400+ students to 50 industry-sponsored projects with a Llama 3.1 and ChromaDB resume parser, reaching 100% fulfillment for every sponsor and student.
- Eliminated 100% of manual score entry with a Zapier workflow that stores student scores from weekly reports in Canvas.
- Tutored undergraduates in Introduction to Statistics: hypothesis testing, probability, and regression fundamentals.

## PROJECTS

**Credit Card Fraud Detection** August 2025
*scikit-learn*, *imbalanced-learn*

- Designed a machine learning pipeline for credit-card fraud classification comparing XGBoost, logistic-regression, and SMOTE resampling models on precision-recall and ROC-AUC, with confusion-matrix-driven, cost-sensitive threshold tuning.

**UpliftBench** February 2026
*Databricks*, *CausalML*, *EconML*, *MLflow*

- Benchmarked S-, T-, X-Learner and Causal Forest uplift models on 13.98M records in PySpark on Databricks, translating CATE rankings into budget-versus-incremental-visit tradeoff curves through policy simulation.

**Financial Forecasting Audit** August 2026
*pandas*, *NumPy*, *SciPy*

- Built a probabilistic financial forecaster and backtested it window by window with Monte Carlo simulation, scored with a rebuilt contest scorer reconciled against every published result.

## SKILLS

**Languages & Cloud:** Python, SQL, R, AWS (SageMaker), Spark, Databricks, Conda
**Machine Learning:** XGBoost, LightGBM, H2O, scikit-learn, PyTorch, Hugging Face Transformers, NLP, Sentiment Analysis
**Modeling & Validation:** Classification, Model Evaluation (ROC-AUC, PR-AUC), Cross-Validation, Backtesting, Time Series, A/B Testing
**MLOps & Data:** MLOps, MLflow, Airflow, Snowflake, PostgreSQL
**Certifications:** AWS Certified Machine Learning Engineer (Associate), Big Data with PySpark (Udemy)

## EDUCATION

**New Jersey Institute of Technology:** MS in Data Science Newark, NJ
**Vellore Institute of Technology:** B.Tech in Computer Science Vellore, India