

Rahul Tamanam

945-308-3656 | rahultamanam24@gmail.com | [linkedin.com/in/rahul-tamanam](https://www.linkedin.com/in/rahul-tamanam) | github.com/rahul-tamanam

SKILLS

Data Engineering: Python, R, SQL, ETL/ELT, Snowflake, Data Pipelines, Data Modeling, PySpark, HDFS, Hive, Impala
Databases and Tools: MySQL, SQLite3, PostgreSQL, MongoDB, Git, Docker, FastAPI
Cloud Platforms: AWS (S3, Redshift, Glue, Lambda, Athena), GCP (BigQuery, Dataflow, Cloud)
Data Analysis and Visualization: Pandas, NumPy, Matplotlib, Seaborn, Power BI, Tableau, Excel
AI/ML & LLM Tools: LangChain, Pinecone, Neo4j, Ollama, Anthropic Claude API, Hugging Face (PEFT, LoRA fine-tuning), vLLM, Axolotl, Weights & Biases, Scikit-learn, XGBoost, spaCy, Groq

EXPERIENCE

AI/ML Engineer & Presence Analytics Specialist July 2026 – TBD

Virtual Village AI

Richardson, TX

- Trained fall detection and presence classification models on time-series mmWave radar point-cloud data using Python and PyTorch, engineering domain-specific features and tracking experiments in Weights & Biases
- Exported trained models to ONNX and validated inference performance on NVIDIA Jetson Orin Nano edge hardware for real-time, on-device deployment
- Designed ground-truth labeling protocols for fall events, presence transitions, and activity counts, and calibrated confidence thresholds to reduce false alarms in a safety-critical detection pipeline
- Tracked model and alert-pipeline performance via Datadog during Alpha Trial deployment across 10 households, documenting findings for stakeholder review

AI/ML Intern

January 2024 – May 2024

HMI Engineering Services

Visakhapatnam, India

- Built a predictive maintenance prototype by cleaning and aggregating time-series sensor data, engineering failure related features, and training classification models using Python, pandas, scikit-learn, and SQL
- Analyzed model predictions to identify high risk equipment and worked with engineers to adjust maintenance thresholds, increasing average equipment uptime by 10%
- Developed failure trend visualizations and sensor drift analyses using Matplotlib and Seaborn to communicate degradation patterns across 6 months of equipment data
- Contributed to an estimated 8% reduction in annual maintenance spend through earlier interventions and fewer repairs

PROJECTS

CometBot | *Python, FastAPI, Neo4j, Pinecone, Typescript* January 2026 - May 2026

- Built CometBot, an AI advising platform for UTD JSOM graduate students, enabling degree planning, career guidance, and skills gap analysis using React, FastAPI, Neo4j, and Pinecone with Claude-powered LLM reasoning
- Designed a hybrid deterministic + RAG architecture where a Neo4j graph enforces academic constraints and degree progression, while Pinecone retrieval grounds LLM outputs, improving reliability of course recommendations
- Architected a deterministic-first evaluation layer that pre-computes all degree requirements before any LLM call, reducing factual hallucination rate from 35% to 0% across 15 tested student profiles

Velox | *React, Node/Express, SQLite, Tailwind CSS, Recharts* April 2026

- Built a macro-aware portfolio engine using FRED (GDP, Fed Funds Rate) to classify market regimes and rebalance an ETF universe monthly, benchmarked against SPY with regime-shaded performance charts
- Ran 1,000-path Monte Carlo simulations via Geometric Brownian Motion, computing p10–p90 fan charts, max drawdown, CAGR, and goal-attainment probability from 5 years of per-ticker price history
- Developed a full-stack investment platform (React, Express, SQLite) with a composite Health Score, live Yahoo Finance integration, LLM-powered assistant, and a portfolio shock simulator

ChurnSense | *Python, SQL, XGBoost, FastAPI, React, LLMs* February 2026 – March 2026

- Built an end-to-end churn pipeline on 30M+ behavioral events, engineering 40+ features in DuckDB (sessionization, 30/60/90-day decay, cohort segmentation) and training an XGBoost model on 860k users with 5-fold CV achieving production-grade accuracy
- Implemented a T-learner uplift model to isolate persuadable users via incremental treatment effects, layered PELT change-point detection for behavioral regime shifts, and deployed SHAP + LLM-generated retention recommendations

EDUCATION

The University of Texas at Dallas May 2026

Master of Science in Business Analytics and Artificial Intelligence

Gayatri Vidya Parishad College of Engineering May 2024

Bachelor of Technology in Information Technology