

CHEBROLU TRILOKESH VENKATA UDAY

Data Scientist | Machine Learning · NLP & LLMs · Model Evaluation · Streaming Data

Andhra Pradesh, India | +91 9063908991 | trilokeshvenkatauday@gmail.com | [GitHub](#) | trilokeshvenkatauday.me

SUMMARY

Data scientist (B.Tech CSE, Data Science, CGPA 8.9) with ~1 year of industry experience at Euron building ML and LLM systems from data collection to deployment. I build models I can explain: time-based validation to avoid leakage, calibrated probabilities, SHAP drivers, honest baseline comparisons, and drift monitoring once live. Experienced with textual data (RAG, NL-to-SQL, LLM-based sentiment and OCR pipelines) and tabular data (XGBoost classifiers on 100K+ row datasets).

EXPERIENCE

Euron

Nov 2025 – Present

Associate Software Engineer

Feb 2026 – Present

- Built **VoteEdge**, an AI election-analytics platform covering **93M+ voters**: Kafka event streaming into a Spark + Delta Lake lakehouse feeding predictive voter scoring and sentiment analysis, deployed multi-region on AWS (primary + DR).
- Designed multi-LLM routing with fallback (LiteLLM across Claude, GPT-4o, Gemini) for OCR result collation and sentiment analysis on text data, with Redis semantic caching to cut LLM cost.
- Designed **GURU**, a multi-agent AI tutoring system (7 engines, ~19 agents) with an assessment engine using confidence-weighted scoring, mastery gates and exam-performance forecasting.

Big Data Engineer Intern

Nov 2025 – Feb 2026

- Built a multi-tenant analytics platform (RAG document Q&A + NL-to-SQL over BigQuery) returning a confidence score and source citations with every answer; monitored latency across RAG (1.2s avg) and BigQuery (850ms avg) layers.
- Designed a multi-agent (LangGraph / CrewAI) NL-to-SQL pipeline for Snowflake with a self-correcting validation loop and human-approved writeback under row-level access controls and audit logging.

PROJECTS

Customer Intelligence Platform: Churn, CLV & Segmentation

[Live demo](#)

- Trained an XGBoost churn model (8K customers, 155K orders, 916K interactions; synthetic) with point-in-time features and quarterly-snapshot validation: **0.86 ROC-AUC, 0.735 PR-AUC, 3.0x top-decile lift** on a held-out period.
- Diagnosed an overconfident model and applied Platt calibration, cutting **Brier score 0.159 → 0.130** with ranking unchanged; benchmarked Random Forest and reported the near-tie (0.736 vs 0.735) rather than hiding it.
- Added per-customer SHAP drivers, PSI drift checks, K-Means segmentation and predictive CLV; shipped with FastAPI, Next.js, Docker and CI (**148 tests**).

Stack: Python, pandas, NumPy, XGBoost, scikit-learn, SHAP, MLflow, FastAPI, PostgreSQL, Kafka, Spark

FraudGuard: Real-Time Fraud Detection

[Live demo](#)

- Built a calibrated fraud classifier and set **approve / review / block** thresholds from the precision–recall trade-off, so uncertain cases go to a human analyst; SHAP explains every flag.
- Scored transactions through Kafka (mTLS) → Spark Structured Streaming using the same code as the API, restart-safe via checkpoints; **124 tests** incl. feature-leakage checks and scheduled drift monitoring.

Stack: scikit-learn, MLflow, SHAP, Kafka, PySpark, FastAPI, PostgreSQL, GitHub Actions

VinoFlow: End-to-End MLOps Pipeline

- Automated the model lifecycle with DVC data/model versioning, MLflow tracking, Jenkins CI/CD retraining and a Dockerised Flask serving API.

TECHNICAL SKILLS

Languages & Core: Python (NumPy, pandas), SQL, Git

Machine Learning: scikit-learn, XGBoost, Random Forest, K-Means, feature engineering, calibration, threshold tuning, SHAP, PyTorch, TensorFlow

NLP & LLMs: RAG, embeddings / pgvector, LangGraph (LangChain), CrewAI, LiteLLM, prompt design, LLM evaluation

Statistics & Evaluation: Probability, linear algebra, hypothesis testing; ROC/PR-AUC, lift, Brier score, calibration, PSI drift

Data, MLOps & Cloud: Kafka, Spark (PySpark), Delta Lake, Databricks, Snowflake, BigQuery, PostgreSQL, MLflow, DVC, Docker, FastAPI, AWS

EDUCATION & PUBLICATION

Bennett University, Greater Noida — B.Tech, CSE (Data Science)

2022 – 2026 | CGPA 8.9/10

Book chapter: *Web Application Security Assessment Techniques: Current Trends and Future Directions — Intelligent Computing and Communication Techniques*, CRC Press (Taylor & Francis), 2025.