

CHEBROLU TRILOKESH VENKATA UDAY

Curriculum Vitae — Data Scientist / ML Engineer

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

PROFILE

Data scientist and ML engineer with a B.Tech in Computer Science (Data Science specialisation, CGPA 8.9, 2026) and about ten months of industry experience at Euron, first as a Big Data Engineer Intern and then as an Associate Software Engineer. My work spans the full model lifecycle: collecting and cleaning data, engineering features, training and evaluating models, and deploying and monitoring them. I care about understanding why a model behaves the way it does, so I validate on later time periods rather than random splits, calibrate probabilities, explain predictions with SHAP, and report results honestly, including when a simpler baseline performs as well. I have worked with textual data (retrieval-augmented generation, NL-to-SQL agents, LLM-based sentiment and OCR pipelines), tabular data (gradient-boosted classifiers on event and transaction data) and streaming systems (Kafka, Spark Structured Streaming).

PROFESSIONAL EXPERIENCE

Euron

Nov 2025 – Aug 2026

Associate Software Engineer

Feb 2026 – Aug 2026

- 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 automatic fallback (LiteLLM across Claude, GPT-4o, Gemini) for OCR result collation and sentiment analysis on unstructured text, with Redis semantic caching to cut LLM cost on repeated queries.
- Designed **GURU**, a multi-agent AI tutoring system (7 engines, ~19 agents) whose assessment engine uses confidence-weighted scoring, mastery gates and exam-performance forecasting to decide what a learner studies next.
- Contributed to **Model Verse**, a generative-AI platform for producing marketing assets.

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 a human-approved writeback flow (validate → approve → merge) under row-level access controls and audit logging.

PROJECTS

Customer Intelligence Platform — Churn, CLV & Segmentation

customerintelligence.trilokeshvenkatauday.me

- Trained an XGBoost churn model on a synthetic dataset of 8K customers, 155K orders and 916K interactions, using point-in-time features and quarterly-snapshot validation to avoid leakage: 0.86 ROC-AUC, 0.735 PR-AUC and 3.0x top-decile lift on a held-out period.
- Found the model was overconfident and applied Platt calibration, reducing Brier score from 0.159 to 0.130 without changing ranking; benchmarked Random Forest and reported the near-tie (0.736 vs 0.735) openly.
- Added per-customer SHAP drivers, PSI drift checks, K-Means segmentation and predictive CLV; shipped with FastAPI, Next.js, Docker and CI covered by 148 automated tests.

Stack: Python, pandas, NumPy, XGBoost, scikit-learn, SHAP, MLflow, FastAPI, PostgreSQL, Next.js, Docker

FraudGuard — Real-Time Fraud Detection

fraudguard.trilokeshvenkatauday.me

- Built a calibrated fraud classifier and set approve / review / block thresholds from the precision-recall trade-off, so uncertain cases are routed to a human analyst; SHAP explains each flag.
- Built a Kafka (Aiven, mTLS) → Spark Structured Streaming pipeline that scores transactions with the same code as the API and writes decisions to PostgreSQL, running as 10-minute micro-batches on GitHub Actions.
- Made it restart-safe (database-backed producer cursor, Spark checkpointing, per-message delivery confirmation) and covered it with 124 automated tests, including feature-leakage checks, plus scheduled drift monitoring.

Stack: Python, scikit-learn, MLflow, SHAP, Kafka, PySpark, FastAPI, PostgreSQL, Next.js, GitHub Actions

VinoFlow — End-to-End MLOps Pipeline

- Automated the ML lifecycle with DVC dataset and model versioning, MLflow experiment tracking, Jenkins CI/CD for continuous training and deployment, and a Dockerised Flask model-serving API.

Alpha-Z — Stock-Market Platform (with a co-founder)

alphaz.trade

- Co-building a live stock-market platform, with hands-on exposure to market data and trader-facing products.

TECHNICAL SKILLS

Programming: Python (NumPy, pandas; clean, tested code), SQL, Git, GitHub Actions

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

NLP & LLMs: Retrieval-augmented generation, embeddings and pgvector, LangGraph (LangChain ecosystem), CrewAI, LiteLLM multi-model routing, NL-to-SQL agents, confidence scoring and citation grounding

Mathematics & Statistics: Probability theory, linear algebra, statistical methods for ML, hypothesis testing

Evaluation & Monitoring: ROC-AUC, PR-AUC, lift, Brier score, calibration curves, PSI drift, MLflow experiment tracking

Data & Infrastructure: Apache Kafka, Spark / PySpark, Spark Structured Streaming, Delta Lake, Databricks, Snowflake, BigQuery, PostgreSQL, MongoDB, Redis, DVC, Docker, FastAPI, AWS, Airflow

EDUCATION

Bennett University, Greater Noida

Sep 2022 – Jun 2026

Bachelor of Technology, Computer Science and Technology (Data Science specialisation) — CGPA 8.9 / 10

PUBLICATION

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

STRENGTHS

- End-to-end ownership: from raw data and features to deployed, tested and monitored models.
- Rigorous, honest evaluation: leakage-safe validation, calibration, and reporting when a baseline ties.
- Comfortable across ML, LLM and big-data stacks, and quick to learn new tools in production settings.

DECLARATION

I hereby declare that the information furnished above is true to the best of my knowledge and belief.

Place: Vijayawada

Chebrolu Trilokesh Venkata Uday

Date: 29 September 2026