

# K. Veeraj Sai Subrahmanyam

veerajsai.me • veeraj.sai@gmail.com • github.com/VeerajSai • linkedin.com/in/veeraj-sai • LeetCode Profile

## Summary

At **Greenko Group**, built a **Ridge-LightGBM ensemble** that cut wind-power scheduling MAE from **7.28 to 3.70 MWh** (49%) across 30+ plants and including a **RAG chatbot** on locally-hosted Ollama LLMs for live operator queries. Research on LLM-driven counterfactual augmentation reduced hate speech classifier bias by **5.02 pp** across 5 classifiers at **7.7 prompts/sec** generation speed, under review at **NeurIPS 2026**. Looking for **AI / Machine learning engineer** roles where production systems and NLP research overlap.

## Education

**Dr. B. R. Ambedkar National Institute of Technology, Jalandhar** **Nov 2022 – May 2026**  
B.Tech in Information Technology

## Publications

**Multimodal Counterfactual Data Augmentation (CAD) for Bias-Aware Hate Content Detection**

K. Veeraj Sai Subrahmanyam *Under Review — NeurIPS 2026*

## Technical Skills

**Languages:** C/C++, Python, Bash, SQL

**ML Frameworks:** PyTorch, TensorFlow, scikit-learn, NumPy, Pandas, Polars, MLFlow, SciPy, AWS Sagemaker

**LLM / GenAI:** HuggingFace Transformers, vLLM, Ollama, RAG Pipelines(Basic), Prompt Engineering, Agentic AI, Qwen3 series

**Tools:** Git, Jupyter, Linux (Arch, Ubuntu), Streamlit, Docker, Langchain, VectorDB(Chroma)

**Cloud / Infra:** Lightning AI (H200 GPU), Kaggle (T4x2 GPU), FastAPI

**Domains:** Machine Learning Engineering, Time-Series Forecasting, NLP, Fairness Evaluation, Open Source.

## Certifications

**AWS Certified AI Practitioner**

*Issued May 3, 2026 – Expires May 3, 2029*

Validation Number: 97d4605ae5914127b8004c15ee85aad6

## Professional Experience

**AI/ML Engineer Intern · Greenko Group**

**Jan 2026 – Feb 2026**

- Built an automated generation-scheduling system for 30+ wind power plants, replacing Excel-based manual workflows with an ML pipeline at 15-minute intervals (96 schedule blocks/day).
- Architected a **Ridge-LightGBM ensemble** that reduced scheduling MAE from **7.28 to 3.70 MWh**—a **49% reduction** against prior manual FSP submissions.
- Provided a **Streamlit interface** for plant operators to monitor live forecasts, compare FSP submissions, and manually override schedule blocks per plant.
- Built a **RAG chatbot** on local Ollama LLMs for natural-language querying of model outputs, feature importances, and plots; integrated into the live Streamlit dashboard and used daily by plant operators.

## Projects

**Hate Speech Bias Mitigation via LLM-Generated Counterfactuals**

*Sep 2025 – Mar 2026*

- Generated **18,000 counterfactual samples** (100% yield) using **Qwen2.5-7B** via a vLLM batch pipeline on Kaggle (T4x2 GPU), with zero cross-class label leakage across 8 protected identity dimensions.
- CAD improved **MiniLM+MLP F1 from 0.863 to 0.956 (+9.6%)** and cut FPR by **75%** (0.237→0.059); validated via OLS ( $F = 9.82, p < 10^{-9}$ ) and Chi-Square ( $\chi^2 = 77.81, p < 10^{-13}$ ).
- Achieved **7.7 CFs/sec** throughput; best text model bootstrap-stable at 95% CI [0.930, 0.960] across 12,000 samples (AUC = 0.979).

**CryptoStats – Real-Time Cryptocurrency Pricing Widget**

*Mar 2025 – Aug 2025*

- Open-source KDE Plasma 6 widget (QML + curl) displaying live Coinbase prices at 5-second refresh, zero external dependencies; published on KDE Store (Pling) for the Linux community.

## Leadership

**Technical Lead — IOTA Club, NITJ**

*Jul 2023 – Present*

Ran ML and Linux workshops for 50+ students; mentored junior members on ML fundamentals and project execution.

## Achievements

- Solved **200+ problems** on **LeetCode** across Data Structures and Algorithms.
- Organized **GDSC HackMol 6.0** as Lead – managed logistics, judging coordination, and execution for 70+ competing teams.
- Ranked **#15,043 in JEE Main** – top 10% among 1.6 million candidates nationwide.