

Aaditya Edupuganti

AI / Machine Learning Engineer

+91-9390257854 | aaditya4401@gmail.com | Gurugram, Haryana, India
linkedin.com/in/aaditya-e | github.com/aaditya4401-dot

AI Engineer with 2 years at Air India applying machine learning to airline operations: an LLM agent platform serving 2,500+ daily users in production, and crew-leave modelling over 525,000 crew-date records. M.Tech in Data Science from IIT Madras, working across tree-based models, time-series forecasting, deep learning and LLMs.

TECHNICAL SKILLS

Languages & Data: Python, SQL, PostgreSQL (star schema), data warehousing, dbt, Prefect, MongoDB, Bash
Modelling: time-series forecasting (Prophet, lag and rolling features), tree-based models (XGBoost, Random Forest), logistic regression, deep learning and transformers (PyTorch, Hugging Face), LLM fine-tuning (QLoRA, DPO/RLHF), anomaly detection, class imbalance, time-based validation, feature engineering
Statistics: hypothesis testing, ANOVA and effect sizes, chi-square, Cramér's V, Mann-Whitney U, Granger causality, stationarity testing, evaluation and metric design
LLM Systems: LangGraph multi-agent orchestration, RAG, structured output, LLM-as-judge evaluation, guardrails, MCP
Production: Azure (OpenAI, AKS), Docker, Kubernetes, microservices, FastAPI, OpenTelemetry, Prometheus, Git, Linux

PROFESSIONAL EXPERIENCE

Air India Limited

AI Engineer

Gurugram, India
Nov 2024 – Present

myDiGi ARIA, LLM agent platform in production

- **Rebuilt Air India's employee assistant end to end**, replacing 12 waterfall dialog flows with **3 LangGraph agents and 14 sub-agents** across 9 microservices on AKS, serving **2,500+ daily active users** at 10.1 model calls per conversation.
- **Owned the crew agent and its sub-agents**, reading roster data and validating fleet eligibility through the CAE Crew Management API at **5.64s P95**, with resumable workflows of up to 14 graph nodes.
- **Kept the LLM outside the security boundary** with a 3-stage pre-inference guard (rule-based checks under 1 ms, then a jailbreak classifier), and shipped the RAG service for policy questions at **1.99s P95**.
- **Instrumented all 9 services** with OpenTelemetry and Prometheus, pinning NLU as the latency bottleneck at **8.3s P95** against a sub-6s target, and cut MongoDB writes from 25 per turn to 1.

Unplanned Crew Leave Prediction, airline operations analytics

Jul 2025 – Oct 2025

- **Modelled unplanned crew leave across 525,000 crew-date records** and 48 base-rank-fleet operational segments, reconstructing historical leave balances per crew member and engineering D-1/D-7 lag features plus a time-since-last-unplanned-leave signal that ranked top by Gini importance at about twice the next feature.
- **Trained logistic regression and random forest classifiers** and judged them on minority-class recall rather than the 85% headline accuracy, since a 9:1 imbalance meant the model missed 77% of unplanned leaves; recommended gradient boosting with lag, rolling and seasonality features as the next phase.
- **Showed holiday-based staffing buffers were not justified:** across 17 hypotheses the holiday effect was insignificant (Cohen's d under 0.08; spike rate 10.16% vs 9.51%, $p=0.33$) while base and fleet explained **28% and 19% of variance**, and flagged DEL/A320 as the highest-risk segment for a spike watchlist.

Kotak Mahindra Bank Limited, Data Science Intern

May 2023 – Jul 2023

Estimated consumer turnover across 30,000 TransUnion CIBIL credit profiles with a tree-based XGBoost regressor at R^2 0.71 on held-out data.

PROJECTS

Vayu: Air Quality Forecasting Pipeline

github.com/aaditya4401-dot

Python, SQL, PostgreSQL, dbt, Prefect, Prophet, statsmodels

2026

- **Forecast 24-hour AQI for 7 Delhi NCR zones** with Prophet over 45 engineered lag, rolling and seasonal features, and **ran 1,344 Granger causality tests** with ADF stationarity checks to isolate 214 significant weather-to-pollution links, on a Prefect pipeline into a PostgreSQL star schema with 10 dbt models and 54 data tests.

RLHF Fine-Tuning of Mistral 7B

github.com/aaditya4401-dot/RLHF

PyTorch, TRL, PEFT, Hugging Face

2026

- **Fine-tuned a 7B transformer with QLoRA SFT and DPO** in 2.5 GPU-hours, training 0.19% of parameters to cut memory from 90 GB to under 16 GB, and built a dual-metric evaluation harness (ROUGE-L against 594 LLM-judge comparisons with Wilson intervals) that exposed judge verbosity bias and train/test overlap in my own split.

EDUCATION

Indian Institute of Technology Madras

Chennai, India

Dual Degree: M.Tech Data Science + B.Tech Aerospace Engineering

2019 – 2024

CGPA: 8.65/10.0 | Coursework: Machine Learning, Statistics, Probability, Optimisation, DSA, DBMS.

Leadership: Team Head, Team Abhyuday, IIT Madras rocketry: 30+ students, 16th of 158, Spaceport America Cup 2023.