

AMBAR NILESH NAIK GURAV

Data Science • Machine Learning • AI Engineering

Goa, India | +91 99754 95001 | ambar0gurav@gmail.com | github.com/ossai | LinkedIn

PROFILE

Data Science undergraduate at **IIT Madras** and **Goa University** with hands-on experience building end-to-end machine learning systems for predictive analytics, NLP, and financial data analysis. Develops production-ready applications using **Python, SQL, FastAPI, Docker, and Streamlit**, combining statistical modeling, feature engineering, and explainable AI to solve real-world business problems. Strong foundation in data analytics, model development, and scalable software engineering, with a focus on fintech and AI-driven decision making.

EDUCATION

IIT Madras (Online) — B.S. in Data Science and Applications Expected 2029 • CGPA 7.5/10

Goa University — Integrated M.Sc. in Data Science Expected 2030 • CGPA 8.7/10

TECHNICAL SKILLS

Programming	Python, SQL, R
Data Analytics	Pandas, NumPy, Data Cleaning, EDA, Statistical Analysis, Feature Engineering, Time-Series Analysis
Machine Learning	Scikit-learn, XGBoost, TensorFlow, Classification, Clustering, SHAP, Cross Validation, Hyperparameter Tuning, Model Evaluation
Visualization	Plotly, Power BI, Tableau, Matplotlib
Backend & Deployment	FastAPI, REST APIs, Streamlit, Docker
Databases & Tools	MySQL, Git, GitHub, Jupyter Notebook, Render

PROJECTS

HireSense AI — Intelligent Resume Screening & Candidate Ranking GitHub

Python • FastAPI • NLP • Docker • Streamlit

- Designed and deployed an end-to-end AI-powered Applicant Tracking System that semantically ranks resumes against job descriptions using transformer embeddings.
- Built scalable FastAPI microservices for resume parsing, candidate scoring, interview question generation, and explainable AI recommendations.
- Containerized with Docker and deployed on Render, enabling real-time resume evaluation via an interactive Streamlit dashboard.
- Engineered a modular architecture supporting bulk resume uploads and transparent recruiter analytics.

Crypto Trader Behavioral Analytics GitHub

Python • Scikit-learn • Plotly • Streamlit

- Built a machine learning pipeline analyzing blockchain trading activity to predict next-day trader profitability from historical transaction behavior.
- Engineered **17 behavioral, sentiment, and rolling-window features**, achieving **62.6% cross-validation accuracy** — roughly **71% above the random baseline**.
- Applied K-Means clustering to discover four distinct trader personas and generated actionable behavioral insights.
- Developed an interactive analytics dashboard for wallet-level prediction, visualization, and model interpretation.

Bank Customer Churn Prediction GitHub

Python • Scikit-learn • XGBoost • SHAP

- Built predictive classification models to identify customers at high risk of churn and support proactive retention strategies.
- Addressed class imbalance using SMOTE, improving recall to **0.634** and reducing false negatives by **58%**.
- Applied SHAP explainability to surface key churn drivers and translate model outputs into business recommendations.

CERTIFICATIONS

Tata Group — GenAI Data Analytics (Forage) • **British Airways** — Data Science (Forage) • **Lloyds Banking Group** — Data Science (Forage)