MOHNISH KUMAR SHANDILYA

Pune, Maharashtra | mohnishshandilya.3000@gmail.com | mohnish4code.github.io
linkedin.com/in/mohnish2004 | github.com/mohnish4code | kaggle.com/mohnishshandilya

Professional Summary

Final-year Electronics and Computer Engineering student who builds end-to-end machine learning systems in computer vision, time series and anomaly detection: from data pipelines, to models evaluated against baselines on chronological and held-out splits, to FastAPI services with React front ends, Docker and CI. Strong foundation in Python, Java, SQL, data structures and algorithms, and object-oriented programming.

Technical Skills

Languages: Python, Java, SQL
Machine Learning & Deep Learning: Scikit-learn, XGBoost, PyTorch, TensorFlow/Keras, Transformers, LSTM, FinBERT, YOLOv8, OpenCV, PaddleOCR, Time-Series Forecasting, Anomaly Detection
Generative AI: RAG, FAISS, Sentence Transformers, LLM APIs (Google Gemini), Prompt Engineering
MLOps & Backend: FastAPI, REST APIs, Docker, GitHub Actions, MLflow, Optuna, Pytest, Postman, Git
Frontend & Data: React, TypeScript, Pandas, NumPy, Drain3, PostgreSQL, Supabase, Streamlit
Core Computer Science: Data Structures and Algorithms, OOP, DBMS, Operating Systems, Computer Networks

Education

MIT ADT University, Pune 2023 – 2027
B.Tech. Electronics & Computer Engineering CGPA: 7.78/10

Projects

System Log Anomaly Detection GitHub
Python, PyTorch, Scikit-learn, Drain3, FastAPI, React, Docker, MLflow

- Built an end-to-end AIOps pipeline over 15.9M raw log lines from the HDFS and BGL datasets; Drain-based parsing reaches 0.998 grouping accuracy on HDFS and 0.969 on BGL.
- Benchmarked 7 detectors (keyword rule, PCA, Isolation Forest, logistic regression, decision tree, DeepLog-style LSTM, causal Transformer) on chronological 60/20/20 splits: best HDFS F1 0.995 (supervised decision tree), best label-free F1 0.793 (PCA); documented why the deep sequence models underperformed on the test period.
- Streaming replay with online parsing reaches F1 0.982 on HDFS (precision 0.968, recall 0.996). FastAPI + React dashboard with explainable alerts; Docker, GitHub Actions CI, 45 tests.

AlphaStream: Multi-Module Stock Prediction GitHub
Python, XGBoost, PyTorch, FinBERT, FastAPI, React, Docker, MLflow, Optuna

- Next-day forecasting system for 10 NSE large caps that combines a forecasting model, technical rules (RSI, SMA, MACD, volatility) and FinBERT news sentiment in a BUY/HOLD/SELL decision engine.
- Compared ARIMA, XGBoost, LSTM and Transformer models, with and without news features, in a leakage-free walk-forward backtest (6 one-year folds, 1,512 trading days per stock, 5-row embargo) against 4 baselines. XGBoost + news is served (net Sharpe 0.47); reported that buy-and-hold (0.65) beat every model.
- FastAPI inference API, React + TypeScript app, Optuna tuning, MLflow tracking, Docker, GitHub Actions CI, 79 tests.

Vehicle Number Plate Detection System GitHub
Python, YOLOv8, PaddleOCR, OpenCV, FastAPI, React, Docker, Google Sheets API

- Custom-trained YOLOv8n plate detector (0.994 mAP@50 on validation) with PaddleOCR (PP-OCRv4) and a plate-aware decoder restricted to valid Indian formats, including Bharat-series and two-row plates.
- On a 68-plate held-out set: 85.3% of plates read exactly (up from 54.4% with EasyOCR), 6.3% character error rate and zero wrong reads auto-accepted; 2.4× faster per image on CPU (4.33 s to 1.81 s).
- FastAPI backend and mobile-friendly React app with phone-camera upload and review before saving; Google Sheets log with duplicate protection and retries.

Other projects: Question Paper Generator (RAG with FAISS, Sentence Transformers, Google Gemini) · Credit Risk Prediction (tuned XGBoost in a scikit-learn Pipeline) · Customer Churn Prediction (TensorFlow/Keras)

Certifications

Kaggle: Introduction to Machine Learning, Intermediate Machine Learning, Feature Engineering · Generative AI Value Added Course, iBase Technologies