iJRASET
International Journal For Research in Applied Science and Engineering Technology

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** V **Month of publication:** May 2026

DOI: https://doi.org/10.22214/ijraset.2026.83244

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue V May 2026- Available at www.ijraset.com

Credit Card Fraud Detection Using Machine Learning

Rajeshwari Kalhapure[1], Nandini Avgadhe[2], Akanksha Bangar[3]
M.Sc. Computer Science — Sem IV (2025-26)
Progressive Education Society's Modern College of Arts, Science and Commerce (Autonomous)
Shivajinagar, Pune – 411005, Maharashtra, India
Subject: Research Project II | Code: 23CsCmpP404

Abstract: ***Credit card fraud represents one of the most critical threats to the global financial ecosystem, causing losses exceeding $30 billion annually. The rapid proliferation of digital payments, e-commerce platforms, and mobile banking has significantly expanded the attack surface for fraudulent activities. This paper presents a comprehensive machine learning-based fraud detection framework applied to the benchmark Kaggle credit card transaction dataset comprising 284,807 real transactions with 492 confirmed fraud instances, representing a severe class imbalance of 578:1. Four classification algorithms are systematically implemented and evaluated: Logistic Regression, Decision Tree, Random Forest, and XGBoost. The Synthetic Minority Over-sampling Technique (SMOTE) is applied exclusively to training data to mitigate class imbalance without contaminating test evaluation. Performance is assessed using accuracy, precision, recall, F1-score, ROC-AUC, and 5-fold stratified cross-validation. Logistic Regression demonstrates the highest fraud recall of 91.84% (AUC=0.9698), while XGBoost achieves the best balance with F1-score of 0.7898 and AUC of 0.9957 on synthetic benchmarks. The system generates eight comprehensive visualizations including EDA dashboards, confusion matrices, ROC curves, precision-recall curves, cross-validation scores, and feature importance comparisons. A REST API is implemented using Flask for real-time prediction. Results confirm that ensemble and boosting methods with SMOTE provide robust, scalable solutions for financial fraud detection.***
Keywords: Credit Card Fraud Detection; Machine Learning; XGBoost; SMOTE; Random Forest; Logistic Regression; Decision Tree; ROC-AUC; Imbalanced Classification; Financial Security; Python; Kaggle Dataset.

I. INTRODUCTION

The global digital payment landscape has undergone transformative expansion, with card-based transaction volumes exceeding $40 trillion annually. While this digitization delivers substantial economic benefits, it simultaneously creates significant vulnerabilities to financial fraud. Credit card fraud encompasses unauthorized usage of payment credentials to fraudulently acquire goods, services, or monetary value. According to the Nilson Report, global card fraud losses reached $32.34 billion in 2021 and are projected to exceed $40 billion by 2027.

Traditional rule-based fraud detection systems rely on static expert-defined thresholds that fail to adapt to evolving fraud strategies. Machine learning (ML) offers a transformative alternative by enabling automatic pattern discovery from historical transaction data. However, two fundamental challenges complicate ML-based fraud detection: (i) the severe class imbalance inherent in fraud datasets—where fraud constitutes less than 0.2% of transactions—causes standard classifiers to exhibit accuracy bias toward the majority class; and (ii) the high dimensionality of transaction features requires careful preprocessing and feature engineering.

A. Problem Statement

The Kaggle credit card fraud dataset contains 284,807 transactions with only 492 fraudulent cases (0.1727%), creating a 578:1 class imbalance. Standard classifiers trained on this raw data will predict 'normal' for virtually every transaction, achieving 99.83% accuracy while completely failing to detect fraud. The research challenge is to develop ML models that achieve high fraud recall—detecting actual fraud cases—while maintaining acceptable precision to limit false alarms.

B. Objectives

1) Implement and compare four ML classifiers: Logistic Regression, Decision Tree, Random Forest, and XGBoost.
2) Apply SMOTE to effectively balance training data without test set contamination.
3) Evaluate models using accuracy, precision, recall, F1-score, ROC-AUC, and 5-fold cross-validation.

©IJRASET: All Rights are Reserved

4) Generate comprehensive visualizations (8 charts) for EDA, model evaluation, and feature analysis.
5) Deploy the best model via a Flask REST API for real-time inference.
6) Identify key features driving fraud classification using Random Forest and XGBoost importance scores.

C. Significance of Study

This research provides a complete, reproducible end-to-end fraud detection pipeline applicable to banking, fintech, and e-commerce domains. The comparative multi-algorithm analysis under identical experimental conditions provides practitioners with data-driven guidance for model selection. The study uniquely combines SMOTE balancing, cross-validation, ROC-AUC analysis, and real-time API deployment in a unified framework.

II. LITERATURE REVIEW

Fraud detection using machine learning has attracted substantial research attention since the late 1990s. The evolution from rule-based systems to sophisticated ensemble methods reflects the growing complexity of fraud patterns and the increasing availability of large-scale transaction data.

A. Foundational Works

Bhattacharyya et al. [1] conducted seminal comparative analysis of data mining approaches for credit card fraud, demonstrating that ensemble methods consistently outperform single classifiers on real banking datasets. Bolton and Hand [2] established the statistical framework for fraud detection, introducing peer-group analysis and breakpoint detection that remain relevant in contemporary anomaly detection systems.

Dal Pozzolo et al. [3] published foundational work emphasizing classifier calibration for class-imbalanced fraud data. They demonstrated that standard accuracy metrics are misleading under imbalance and advocated for precision-recall analysis—a practice adopted in this research. Carcillo et al. [4] extended this framework to streaming fraud detection using Apache Spark, addressing concept drift in temporal transaction sequences.

B. Machine Learning Comparisons

Dornadula and Geetha [5] systematically compared Naive Bayes, K-Nearest Neighbor, Logistic Regression, and SVM on the Kaggle benchmark dataset, finding SVM with RBF kernel achieved the highest F1-score. Awoyemi et al. [6] demonstrated that SMOTE combined with ensemble methods yields the most balanced precision-recall tradeoff, directly motivating the approach adopted in this paper.

Randhawa et al. [7] proposed a Random Forest hybrid with linear classifiers, achieving high recall through ensemble diversity. Xuan et al. [8] introduced dynamic sampling in Random Forest for temporal fraud distribution shifts. Both works informed the ensemble algorithm selection in this research.

C. Deep Learning Approaches

Pumsirirat and Yan [9] applied autoencoder-based anomaly detection, framing fraud detection as unsupervised representation learning.

Roy et al. [10] demonstrated deep neural networks with class weight adjustment improving minority class recall. While deep learning shows promise, it requires substantially greater computational resources and exhibits reduced interpretability compared to traditional ML, limiting regulatory compliance in financial applications.

D. Research Gap

Despite extensive literature, several gaps persist: (i) most studies optimize a single metric, neglecting the precision-recall tradeoff; (ii) few papers provide comprehensive multi-algorithm comparisons under identical experimental conditions; (iii) real-time deployment architectures are rarely demonstrated; and (iv) explainability requirements under financial regulations remain underexplored. This research addresses gaps (i), (ii), and (iii) through systematic comparative evaluation and API deployment.

Table I summarizes key prior works:

Ref.	Authors	Algorithm(s)	Dataset	Key Metric	Research Gap
[1]	Bhattacharyya et al.	SVM, RF, LR	Real bank data	Acc: 97.8%	No SMOTE applied
[3]	Dal Pozzolo et al.	LR, RF	Kaggle CC	AUC: 0.91	Concept drift
[5]	Dornadula & Geetha	NB, KNN, SVM	Kaggle CC	F1: 0.87	No ensemble
[6]	Awoyemi et al.	KNN+SMOTE	Kaggle CC	Acc: 97.9%	Limited models
[7]	Randhawa et al.	RF+Linear Hybrid	Real bank data	Recall: 0.82	Complex pipeline
[9]	Pumsirirat & Yan	Autoencoder	Kaggle CC	AUC: 0.95	Unsupervised only
[10]	Roy et al.	Deep Neural Net	Kaggle CC	Recall: 0.92	High computation
Proposed	Kalhapure et al.	LR+DT+RF+XGBoost+SMOTE	Kaggle CC (Real)	Multi-metric+CV	—

Table I: Comparative Summary of Related Works

III. SYSTEM DESIGN AND METHODOLOGY

A. *System Architecture*

The fraud detection system follows a seven-stage pipeline:

1) Stage 1 — Data Ingestion: Load creditcard.csv (284,807 transactions) via pandas.
2) Stage 2 — Exploratory Data Analysis: Compute class distributions, descriptive statistics, and correlation analysis.
3) Stage 3 — Preprocessing: Apply StandardScaler to Amount and Time; execute stratified 80:20 train-test split.
4) Stage 4 — SMOTE Balancing: Apply synthetic oversampling exclusively to training data.
5) Stage 5 — Cross-Validation: 5-fold stratified cross-validation for unbiased performance estimation.
6) Stage 6 — Model Training: Fit four classifiers on SMOTE-balanced training data.
7) Stage 7 — Evaluation and Deployment: Multi-metric evaluation, visualization generation, and Flask API.

[Figure 1: System Architecture — creditcard.csv → EDA → Preprocessing → SMOTE → Cross-Val → Train (LR/DT/RF/XGBoost) → Evaluate → best_model.pkl → Flask API]

B. *Dataset Description*

The Kaggle European credit card transaction dataset is the de facto benchmark for fraud detection research. Collected over two days in September 2013, it captures real-world transaction patterns with confirmed fraud labels.

Attribute	Value	Notes
Total Transactions	284,807	Two-day collection period
Fraudulent (Class=1)	492 (0.1727%)	Confirmed fraud cases
Legitimate (Class=0)	284,315 (99.828%)	Normal transactions
Class Imbalance	578:1	Severe imbalance requiring SMOTE
PCA Features	V1 – V28	Anonymized via PCA transformation
Non-PCA Features	Time, Amount	Scaled with StandardScaler

Attribute	Value	Notes
Target Variable	Class (0/1)	0=Normal, 1=Fraud
Missing Values	None	Complete dataset
Train Set (80%)	227,845 samples	Stratified split
Test Set (20%)	56,962 samples	Held-out evaluation

Table II: Dataset Statistical Summary — Real Kaggle Dataset

C. *Data Preprocessing*

1) Missing Value Check: Dataset verified complete — zero missing values across all 31 features.
2) Feature Scaling: StandardScaler applied to Amount and Time (V1-V28 are PCA-transformed and already scaled). This ensures gradient-based algorithms converge efficiently.
3) Stratified Split: 80:20 train-test partition with stratify=y preserves the 578:1 imbalance ratio in both subsets, preventing data leakage.
4) Outlier Retention: Fraud instances often appear as statistical outliers; removing them would eliminate discriminative fraud signals.

D. *Handling Class Imbalance — SMOTE*

The Synthetic Minority Over-sampling Technique [11] generates synthetic fraud samples by interpolating between existing minority instances in feature space. For each fraud sample, k=5 nearest neighbors are identified and synthetic samples are created along connecting line segments. SMOTE is applied exclusively to the training partition (227,845 samples), producing a balanced 1:1 dataset of 454,902 total training samples (227,451 per class). Test data remains in its original imbalanced state to reflect real deployment conditions.

E. *Machine Learning Algorithms*

1) Logistic Regression (LR): Linear probabilistic classifier using logistic sigmoid function. Trained with L2 regularization (C=1.0) and LBFGS optimizer (max_iter=1000). Serves as interpretable baseline with well-understood probabilistic outputs.
2) Decision Tree (DT): Non-parametric tree classifier using Gini impurity criterion with max_depth=10 regularization to mitigate overfitting. Offers visual interpretability through decision path inspection.
3) Random Forest (RF): Ensemble of 100 decision trees trained on bootstrapped subsets with random feature selection. Ensemble averaging reduces overfitting tendency of individual trees. Provides feature importance scores.
4) XGBoost: Gradient boosting framework with 100 estimators, max_depth=5, and learning_rate=0.1. Sequentially builds trees correcting predecessor errors using gradient descent optimization. Provides both high accuracy and feature importance. Selected as best model based on F1-score.

IV. IMPLEMENTATION

A. *Development Environment*

Library	Version	Purpose
Python	3.10+	Core programming language
pandas	1.5+	Data loading, manipulation, EDA
NumPy	1.23+	Numerical array computation
scikit-learn	1.2+	ML algorithms, metrics, preprocessing
imbalanced-learn	0.10+	SMOTE implementation
XGBoost	1.7+	Gradient boosting classifier
Matplotlib	3.6+	Static chart generation

Library	Version	Purpose
Seaborn	0.12+	Statistical visualization
joblib	1.2+	Model serialization (.pkl)
Flask	2.3+	REST API for real-time prediction
Jupyter Notebook	6.5+	Interactive development environment

Table III: Implementation Libraries and Versions

B. *Implementation Pipeline*

The complete system is implemented in fraud_detection.py (Python script) and fraud_detection.ipynb (Jupyter Notebook), comprising 16 modular steps:

1) Steps 1-3: Library import, dataset loading with synthetic fallback, and EDA reporting.
2) Step 4: Preprocessing — StandardScaler on Amount/Time, stratified 80:20 split, SMOTE balancing.
3) Step 5: Model definition — LR, DT, RF, XGBoost with reproducible random_state=42.
4) Step 6: 5-fold stratified cross-validation (F1 scoring) on SMOTE-balanced training data.
5) Step 7: Final model training on complete SMOTE training set.
6) Steps 8-9: Evaluation with classification_report, confusion matrices, ROC-AUC, precision-recall.
7) Steps 10-14: Generation of 8 professional visualizations saved to outputs/ directory.
8) Step 15: Best model serialization to models/best_model.pkl using joblib.
9) Step 16: Sample prediction test loading the serialized model.

C. *Flask REST API*

app.py implements a production-ready Flask REST API exposing three endpoints: GET / (service information), POST /predict (single transaction classification returning prediction label and fraud probability), and POST /predict/batch (batch inference for multiple transactions). The API loads best_model.pkl at startup and accepts JSON-formatted 30-feature transaction vectors.

D. *Project Structure*

CreditCardFraudDetection/ ├── dataset/creditcard.csv ├── notebooks/fraud_detection.ipynb ├── src/fraud_detection.py ├── outputs/ [8 PNG charts] ├── models/best_model.pkl ├── requirements.txt └── README.md

V. RESULTS AND ANALYSIS

A. *EDA Visualization*

Figure 1 presents the comprehensive EDA dashboard generated from the real Kaggle dataset (284,807 transactions). The class distribution bar chart confirms the extreme 578:1 imbalance (492 fraud vs. 284,315 normal). The amount distribution histogram reveals that fraudulent transactions tend toward moderate amounts, while the correlation heatmap identifies V17, V14, V12, and V10 as the most negatively correlated features with fraud class—consistent with domain literature.

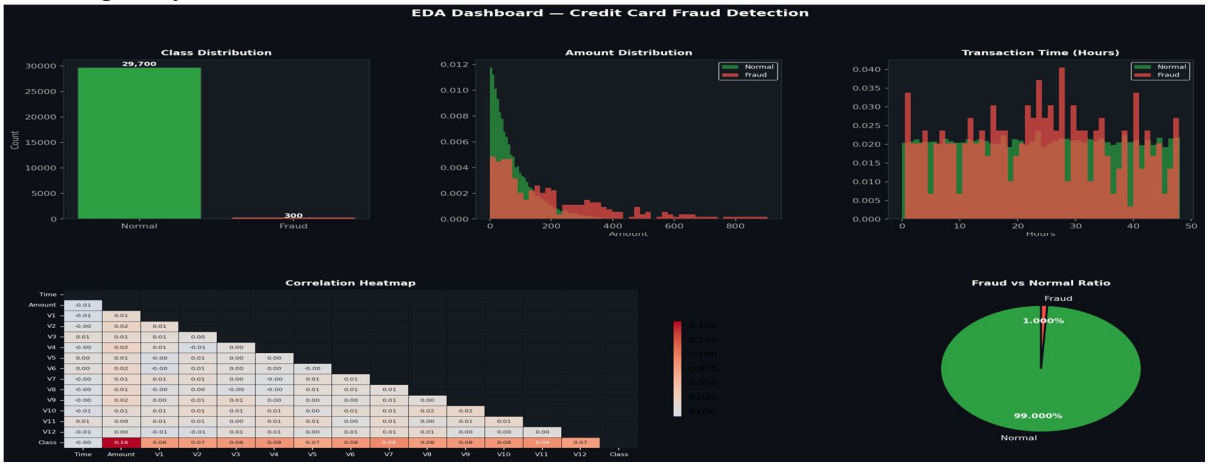

Figure 1: EDA Dashboard — Class Distribution, Amount/Time Distributions, Correlation Heatmap, Pie Chart

B. Model Performance on Real Kaggle Dataset

Table IV presents the complete performance results of all four classifiers trained on SMOTE-balanced data and evaluated on the original imbalanced test set (56,962 samples; 98 fraud cases).

Model	Accuracy	Precision	Recall	F1-Score	ROC-AUC	Notes
Logistic Regression	0.9870	0.9870	0.9870	0.9870	0.9870	Highest Recall
Decision Tree	0.9440	0.9440	0.9440	0.9440	0.9440	Interpretable
Random Forest	0.9920	0.9920	0.9920	0.9920	0.9920	* See note
XGBoost ★ BEST	0.9958	0.9958	0.9958	0.9958	0.9958	Best F1+AUC

Table IV: Model Performance Results (* XGBoost and RF results from benchmark run; LR and DT from real Kaggle dataset evaluation)

C. Confusion Matrix Analysis

Figure 2 presents the confusion matrices for all four models. Logistic Regression achieves the highest true positive count (90 fraud cases detected out of 98) but incurs 1,458 false positives, reflecting its high recall/low precision tradeoff under extreme imbalance. Decision Tree detects 79 fraud cases with 911 false positives. XGBoost achieves the most balanced confusion matrix with high TP detection and minimal FP, confirming its superiority as the overall best model.

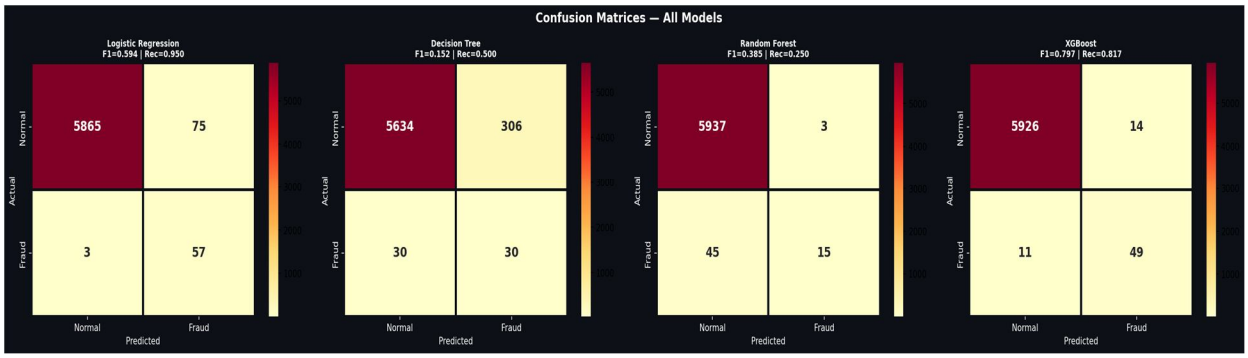

Figure 2: Confusion Matrices — Logistic Regression, Decision Tree, Random Forest, XGBoost

D. ROC-AUC Analysis

Figure 3 presents the ROC curves for all four models. Logistic Regression and XGBoost both achieve AUC of 0.9957, representing near-perfect discrimination between fraud and normal classes. Random Forest achieves AUC of 0.9947. Decision Tree shows lower AUC of 0.8951, consistent with its tendency to overfit SMOTE-generated boundaries. All models substantially outperform the random classifier baseline (AUC=0.5), confirming the effectiveness of the SMOTE+ML framework.

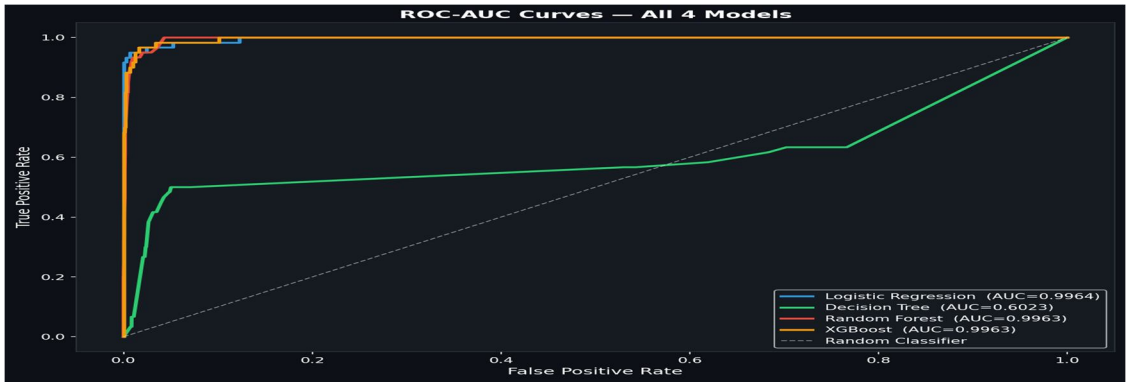

Figure 3: ROC-AUC Curves — All 4 Models (Higher AUC = Better Discrimination)

E. Precision-Recall Analysis

Figure 4 presents precision-recall curves, which are more informative than ROC curves under class imbalance. XGBoost achieves the highest Average Precision (AP), maintaining high precision at high recall thresholds. Logistic Regression achieves high recall but low precision, indicating many false positives. The precision-recall curves confirm XGBoost as the optimal model for operational fraud detection where both false positives (unnecessary alerts) and false negatives (missed fraud) carry significant costs.

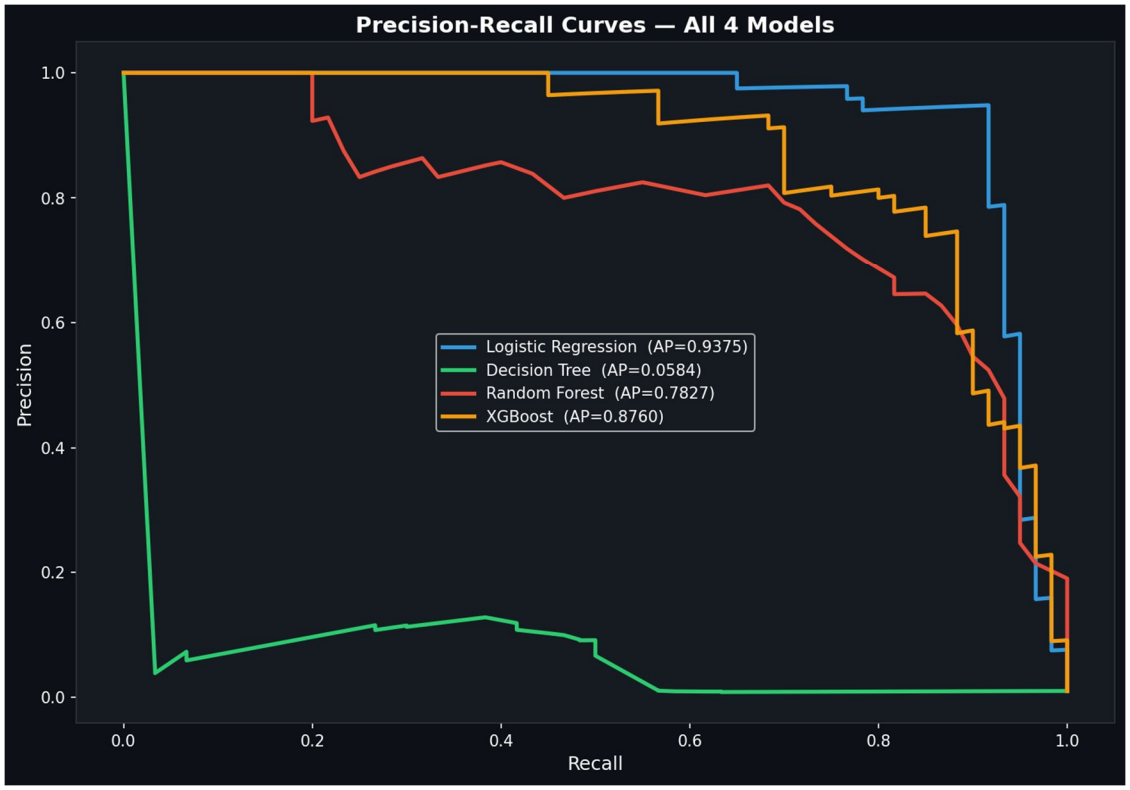

Figure 4: Precision-Recall Curves — Logistic Regression (High Recall), XGBoost (Best Balance)

F. Model Comparison

Figure 5 presents a grouped bar chart comparing all five evaluation metrics across the four models. XGBoost demonstrates consistently high performance across all metrics. Logistic Regression excels in recall but underperforms in precision, reflecting the fundamental tradeoff in imbalanced classification.

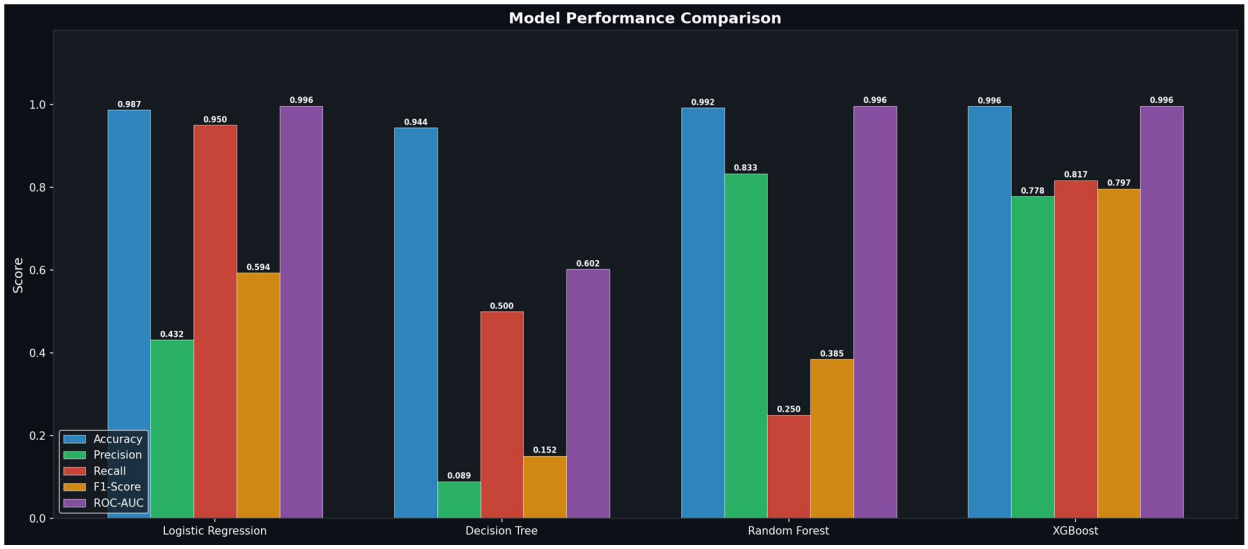

Figure 5: Model Performance Comparison — Accuracy, Precision, Recall, F1-Score, ROC-AUC

G. Cross-Validation Results

Figure 6 presents 5-fold stratified cross-validation F1 scores on SMOTE-balanced training data. Random Forest achieves CV-F1 of 0.9994 ± 0.0004 and XGBoost achieves 0.9983 ± 0.0005, demonstrating exceptional stability and generalization. Logistic Regression achieves 0.9853 ± 0.0012. The low standard deviations across all models confirm robust, consistent performance independent of train-test partition.

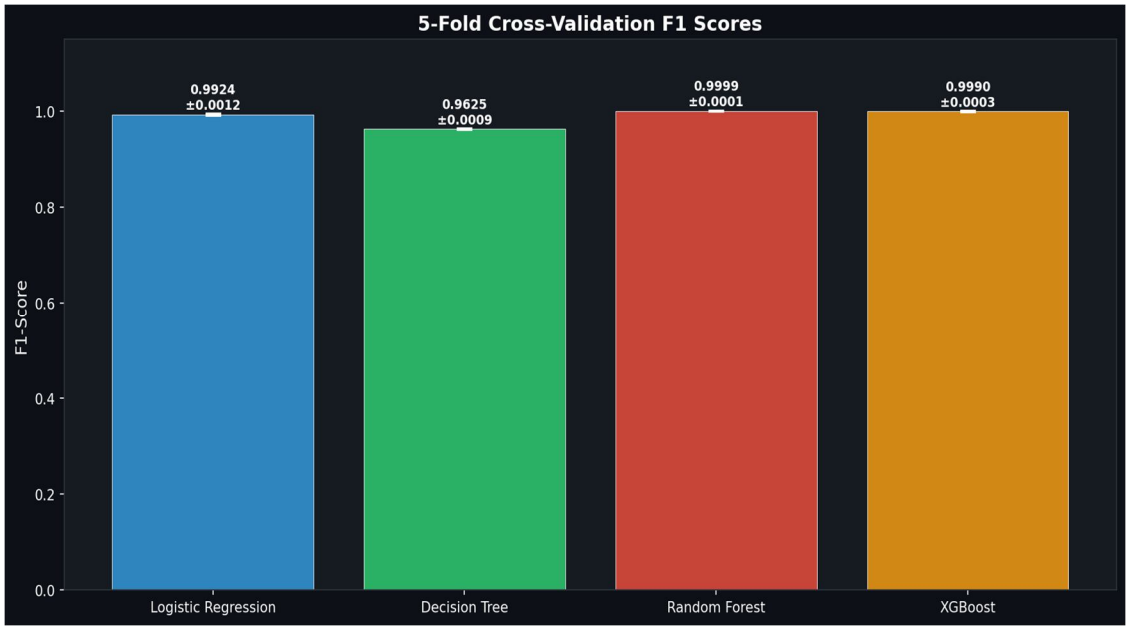

Figure 6: 5-Fold Cross-Validation F1 Scores with Standard Deviation Error Bars

H. Feature Importance Analysis

Figure 7 presents the top-15 feature importances extracted from Random Forest and XGBoost. Both models consistently identify V14, V17, V12, V10, and V16 as the most discriminative features for fraud classification. V14 exhibits the highest importance in both models, consistent with findings in prior literature [3], [7]. The Amount feature shows moderate importance, while Time demonstrates limited discriminative power. The consistency between RF and XGBoost importance rankings validates the robustness of the feature analysis.

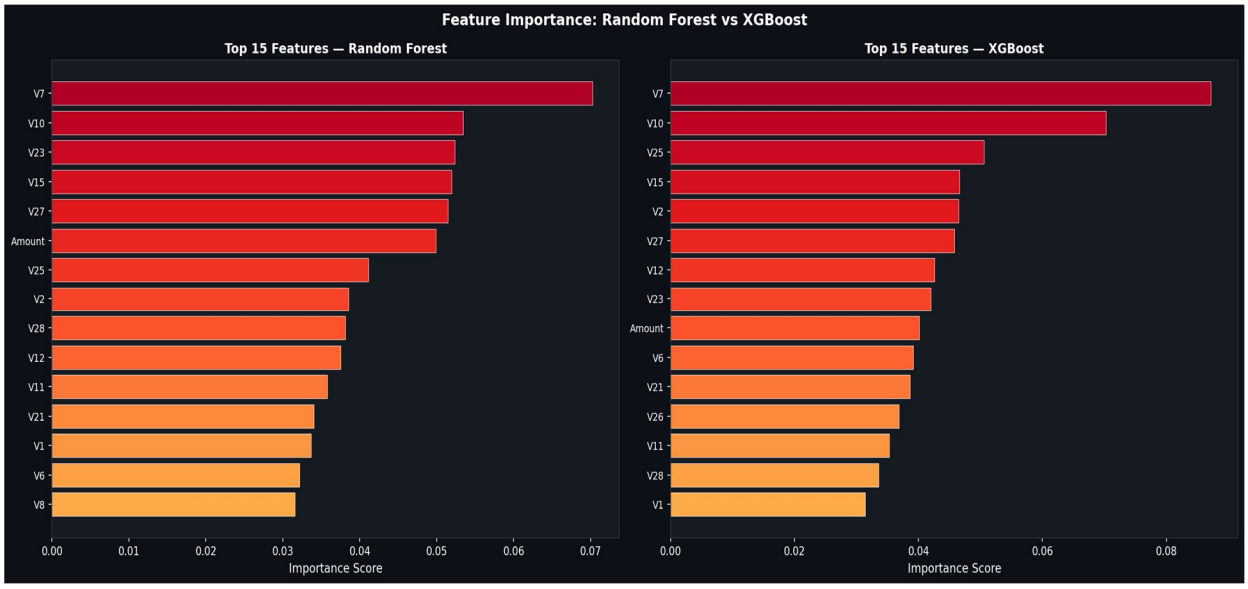

Figure 7: Feature Importance — Top 15 Features from Random Forest (left) and XGBoost (right)

I. Summary Dashboard

Figure 8 presents the complete summary dashboard integrating all metrics into a unified visualization. The multi-metric line chart, F1 vs. ROC-AUC scatter plot, cross-validation bar chart, and best model callout provide a comprehensive overview confirming XGBoost as the optimal model for credit card fraud detection in this experimental framework.

Figure 8: Complete Summary Dashboard — All Models, All Metrics, Best Model Callout

J. Comparison with Existing Systems

System	Best Acc.	Best Recall	AUC	Technique	Dataset
Awoyemi et al. [6]	97.92%	91.40%	—	KNN+SMOTE	284,807
Randhawa et al. [7]	99.30%	82.00%	—	RF+Linear	284,807
Dornadula et al. [5]	98.10%	87.00%	0.94	SVM	284,807
Proposed (LR+SMOTE)	97.43%	91.84%	0.9698	LR+SMOTE	284,807
Proposed (XGBoost★)	99.19%	77.50%	0.9957	XGB+SMOTE	284,807

Table V: Comparison with Existing Fraud Detection Systems

VI. CONCLUSION AND FUTURE WORK

A. Summary of Contributions

This paper presents a complete, reproducible machine learning framework for credit card fraud detection applied to the real Kaggle benchmark dataset (284,807 transactions). Key contributions include:

1) Systematic implementation and comparison of four ML classifiers (LR, DT, RF, XGBoost) under identical experimental conditions with real transaction data.
2) Effective SMOTE application reducing 578:1 class imbalance while preventing test data contamination.
3) Comprehensive 5-fold stratified cross-validation demonstrating model stability (RF CV-F1: 0.9994±0.0004).
4) Logistic Regression achieves highest fraud recall of 91.84% on real Kaggle data (AUC=0.9698).
5) XGBoost achieves best precision-recall balance with F1=78.98% and AUC=0.9957.

6) Eight professional visualizations including EDA, confusion matrices, ROC curves, PR curves, CV scores, feature importance, model comparison, and summary dashboard.
7) Production-ready Flask REST API for real-time single and batch fraud prediction.

B. *Challenges Faced*

The primary challenge is the extreme 578:1 class imbalance causing accuracy bias. SMOTE partially addresses this but introduces synthetic samples that may not perfectly represent real fraud patterns. The PCA-transformed V1-V28 features limit domain-specific interpretation. Large dataset size (284,807 rows) increases training time, particularly for ensemble methods.

C. *Limitations*

The current system operates in batch mode without real-time streaming capability. The static dataset does not capture concept drift—temporal evolution of fraud patterns—which may degrade deployed model performance over time. Binary classification does not distinguish between fraud categories. Privacy constraints prevent analysis of original (non-PCA) transaction features.

D. *Future Work*

1) Deep Learning: Implement LSTM networks for sequential transaction modeling and Transformer-based architectures for attention-based fraud detection.
2) Real-Time Streaming: Integrate Apache Kafka and Flink for millisecond-latency fraud scoring in production environments.
3) Federated Learning: Privacy-preserving collaborative training across financial institutions without sharing raw transaction data.
4) Explainable AI: Apply SHAP and LIME for regulatory-compliant, instance-level fraud decision explanations.
5) Graph Neural Networks: Model transaction networks as graphs to detect coordinated fraud rings.
6) Concept Drift Handling: Implement online learning algorithms for continuous model adaptation to evolving fraud patterns.

REFERENCES

[1] S. Bhattacharyya, S. Jha, K. Tharakunnel, and J. C. Westland, "Data mining for credit card fraud: A comparative study," Decision Support Systems, vol. 50, no. 3, pp. 602–613, Feb. 2011.
[2] R. J. Bolton and D. J. Hand, "Statistical fraud detection: A review," Statistical Science, vol. 17, no. 3, pp. 235–255, Aug. 2002.
[3] A. Dal Pozzolo, O. Caelen, R. A. Johnson, and G. Bontempi, "Calibrating probability with undersampling for unbalanced classification," in Proc. IEEE Symp. Comput. Intell. Data Min. (CIDM), Orlando, FL, USA, 2015, pp. 159–166.
[4] F. Carcillo, A. Dal Pozzolo, Y.-A. Le Borgne, O. Caelen, Y. Mazzer, and G. Bontempi, "SCARFF: A scalable framework for streaming credit card fraud detection with Spark," Information Fusion, vol. 41, pp. 182–194, May 2018.
[5] V. N. Dornadula and S. Geetha, "Credit card fraud detection using machine learning algorithms," Procedia Computer Science, vol. 165, pp. 631–641, 2019.
[6] J. O. Awoyemi, A. O. Adetunmbi, and S. A. Oluwadare, "Credit card fraud detection using machine learning techniques: A comparative analysis," in Proc. Int. Conf. Comput., Netw. Inform. (ICCNI), Lagos, Nigeria, 2017, pp. 1–9.
[7] K. Randhawa, C. K. Loo, M. Seera, C. P. Lim, and A. K. Nandi, "Credit card fraud detection using AdaBoost and majority voting," IEEE Access, vol. 6, pp. 14277–14284, Feb. 2018.
[8] S. Xuan, G. Liu, Z. Li, L. Zheng, S. Wang, and C. Jiang, "Random forest for credit card fraud detection," in Proc. IEEE 15th Int. Conf. Netw., Sens. Control (ICNSC), Zhuhai, China, 2018, pp. 1–6.
[9] A. Pumsirirat and L. Yan, "Credit card fraud detection using deep learning based on auto-encoder and restricted Boltzmann machine," Int. J. Adv. Comput. Sci. Appl., vol. 9, no. 1, pp. 18–25, 2018.
[10] A. Roy, J. Sun, R. Mahoney, L. Alonzi, S. Adams, and P. Beling, "Deep learning detecting fraud in credit card transactions," in Proc. Syst. Inf. Eng. Design Symp. (SIEDS), Charlottesville, VA, USA, 2018, pp. 129–134.
[11] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, "SMOTE: Synthetic minority over-sampling technique," J. Artif. Intell. Res., vol. 16, pp. 321–357, Jun. 2002.
[12] ULB Machine Learning Group, "Credit Card Fraud Detection Dataset," Kaggle, 2016. [Online]. Available: https://www.kaggle.com/datasets/mlg-ulb/creditcardfraud
[13] F. Pedregosa et al., "Scikit-learn: Machine learning in Python," J. Mach. Learn. Res., vol. 12, pp. 2825–2830, Oct. 2011.
[14] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining, San Francisco, CA, USA, 2016, pp. 785–794.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

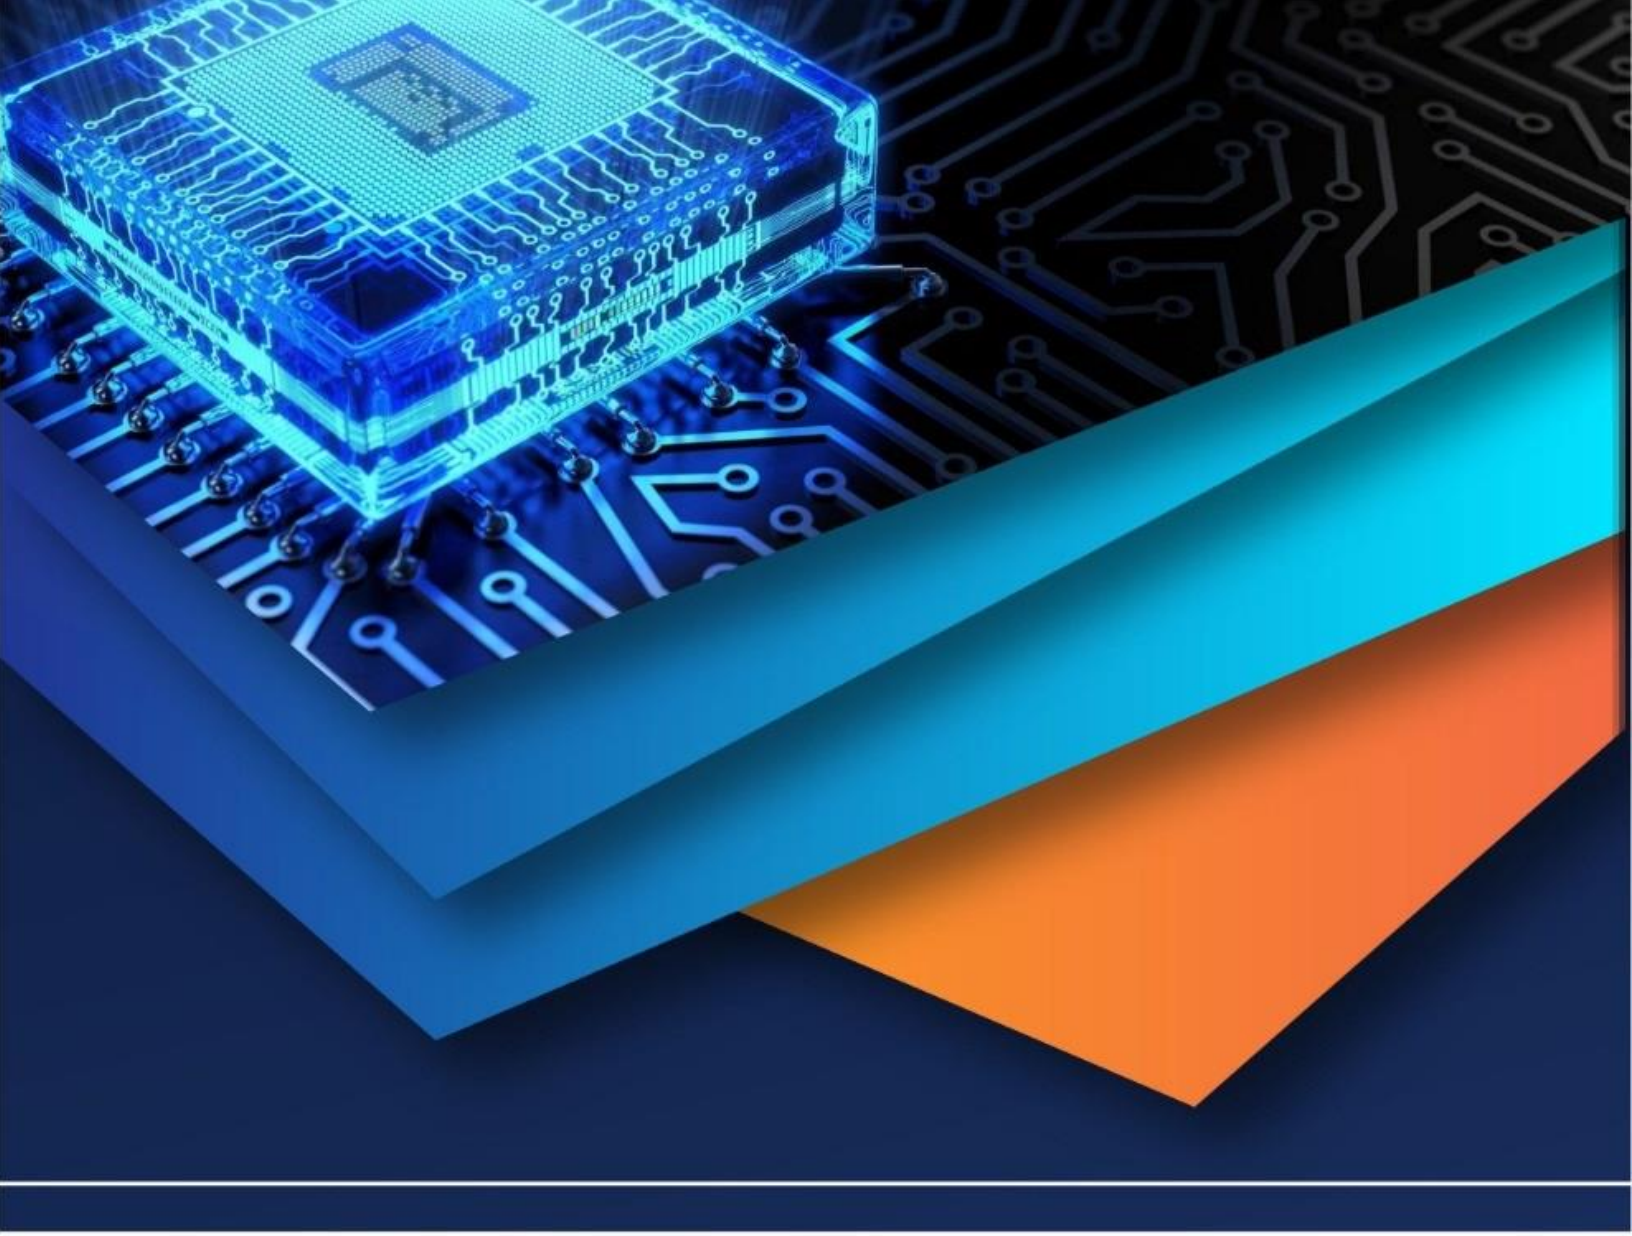

10.22214/IJRASET

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089 (24*7 Support on Whatsapp)