IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

Credit Card Fraud Detection Using Machine Learning

Yagnesh Mittal[1], Syed Abdullah Affan[2], Adarsh Shah[3] , K.K. Sharma[4]

[123]B. Tech. Student, [4]Professor

[1234]Department of Information Technology

Shri G. S. Institute of Technology and Science, Indore, M.P., India

Abstract:

The rapid growth of digital payment systems has led to a corresponding increase in credit card fraud, creating serious financial and security concerns for institutions and consumers worldwide. Timely and accurate identification of fraudulent transactions has therefore become essential. This study presents a machine learning–based approach for credit card fraud detection, aimed at improving the reliability of real-time transaction monitoring.

The proposed framework leverages historical transaction data to learn behavioral patterns and detect anomalies without relying on predefined rules. Emphasis is placed on data preprocessing techniques to manage class imbalance, along with feature engineering methods to enhance predictive performance. Several supervised learning models, including Logistic Regression, Decision Trees, Random Forests, and XG-Boost, are implemented and evaluated for classification accuracy.

The overall objective of this work is to effectively distinguish fraudulent transactions from legitimate ones, thereby strengthening transaction security and minimizing financial losses.

Keywords: Credit Card Fraud Detection, Machine Learning, Anomaly Detection, Imbalanced Data, Supervised Learning, Transaction Security

I. INTRODUCTION

The continuous rise in digital and card-based financial transactions has significantly increased exposure to credit card fraud, presenting a serious challenge for modern financial systems. As transaction volumes grow and payment platforms become more complex, fraudulent activities have evolved in sophistication, making timely and accurate detection increasingly difficult. Existing fraud detection mechanisms, which largely rely on static rules and manual verification, often struggle to identify complex fraud patterns and fail to operate effectively in real-time environments. This limitation leads to increased financial losses, operational inefficiencies, and reduced customer trust.

The primary objective of a credit card fraud detection system is to accurately differentiate between legitimate and fraudulent transactions while maintaining minimal disruption to genuine users. Achieving high detection accuracy must be balanced with the reduction of false positives, as incorrect transaction blocking can negatively impact customer experience. Additionally, fraud detection systems must be capable of responding in real time to prevent unauthorized transactions before financial damage occurs.

Machine learning techniques provide an effective solution to these challenges by enabling automated analysis of large-scale transaction data and identification of anomalous patterns. By learning from historical data, machine learning models can adapt to evolving fraud strategies without the need for constant manual rule updates. Python-based analytical frameworks further support this approach by offering scalable tools for data preprocessing, feature extraction, and model development.

Another key objective is system adaptability and scalability. As fraud tactics continuously change, the detection model must update its learning from new transaction data to remain effective over time. Furthermore, the system should securely handle high transaction volumes while ensuring consistent performance. By addressing these objectives, machine learning-driven credit card fraud detection systems offer a robust and efficient solution to the growing challenges of financial fraud in digital payment ecosystems.

II. LITERATURE REVIEW

Credit card fraud detection has evolved from static, rule-based mechanisms to data-driven machine learning frameworks capable of modeling complex and non-linear transactional behavior. Early research demonstrated that data mining techniques outperform handcrafted rules by learning fraud signatures directly from historical transactions, enabling improved detection in dynamic environments by [1] (Bhattacharyya, 2011). Subsequent survey studies confirmed that supervised learning models provide superior adaptability to emerging fraud patterns compared to traditional heuristic systems by [2][5] (Chandran).

Foundational work in data mining emphasizes the importance of feature engineering, noise handling, and rigorous validation protocols for building reliable predictive systems by [3][8](Han, 2012). In the context of fraud detection, classification and ensemble learning methods have been shown to offer higher robustness in noisy, high-dimensional financial data. A comprehensive review by [4] (Kotsiantis, 2006) highlights that combining classifiers can significantly improve generalization performance, which is particularly beneficial for highly imbalanced problems such as fraud detection.

Among classical machine learning approaches, Logistic Regression remains a strong baseline due to its interpretability and stable calibration properties, which are desirable in regulated financial environments. Empirical findings suggest that logistic models are computationally efficient and suitable for real-time deployment; however, their linear decision boundaries limit performance when fraud patterns exhibit complex feature interactions [7] (Waleed).

Tree-based models provide improved expressiveness by capturing non-linear relationships among transaction attributes such as amount, temporal behavior, and user spending patterns. Random Forest, as an ensemble of decision trees, reduces overfitting and improves detection stability. Empirical evaluations indicate that Random Forest classifiers achieve higher recall at controlled false-positive rates compared to single-tree models, making them suitable for operational fraud detection systems [1] (Bhattacharyya, 2011)[2] (Chandran).

Gradient boosting techniques, particularly XGBoost, further enhance performance by optimizing regularized objective functions over complex feature interactions. Studies report that boosted ensembles achieve superior precision–recall trade-offs under severe class imbalance when appropriately tuned and (Kotsiantis, 2006) calibrated [5] (Monika. R. & Shukla, 2019). Support Vector Machines have also been explored for fraud detection and demonstrate strong classification performance on smaller or well-structured datasets; however, scalability constraints and kernel selection challenges may limit their applicability in large-scale, real-time systems [6] [[7][13](Shahrukh).

Class imbalance remains a fundamental challenge in fraud detection, as fraudulent transactions represent only a small fraction of all records. Prior research emphasizes that accuracy is an inadequate metric in such scenarios and recommends the use of precision, recall, F1-score, and PR-AUC for reliable evaluation [1][2] (Chandran) [3] (Han, 2012). Cost-sensitive learning and threshold optimization are also advocated to align model decisions with operational risk constraints.

Overall, the literature supports the adoption of ensemble learning methods, particularly Random Forest and XGBoost, evaluated using precision–recall-oriented metrics and deployed with interpretability considerations. These approaches provide a balanced trade-off between detection performance, scalability, and governance, motivating their selection for real-time credit card fraud detection systems[6] [1] (Bhattacharyya, 2011)[5] (Monika. R. & Shukla, 2019).

III. METHODOLOGY

Figure 3.1 Illustrates the sequence diagram in a visual manner, enabling you both to document and validate your logic, and commonly used for both analysis and design purpose.

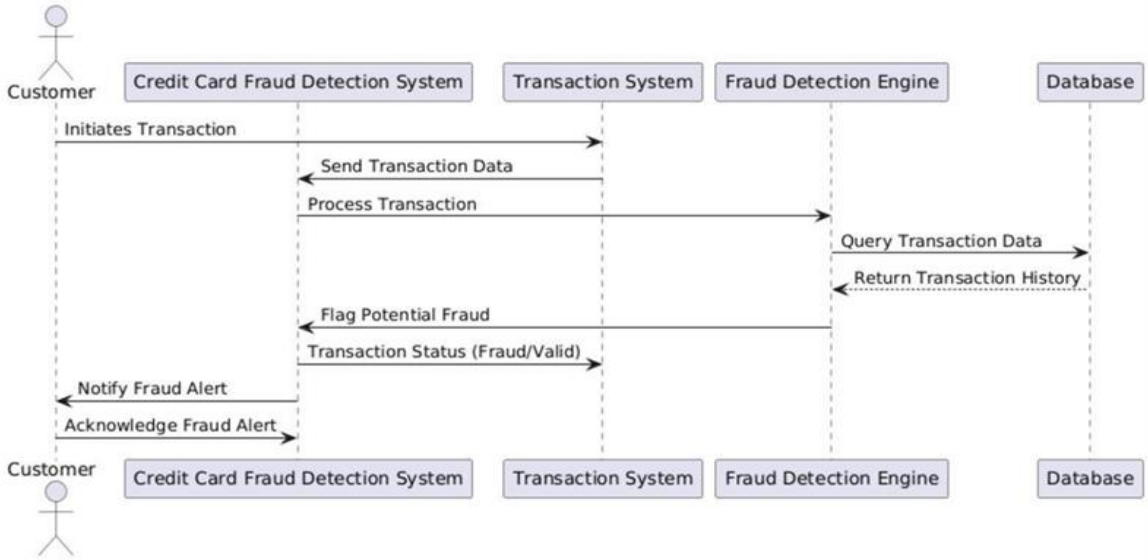

Sequence diagram illustrating system logic and interactions

Figure 3.2 illustrates the architecture of the proposed framework, outlining the complete workflow for credit card fraud detection. The effectiveness of the system is subsequently evaluated using standard performance metrics to assess model performance.

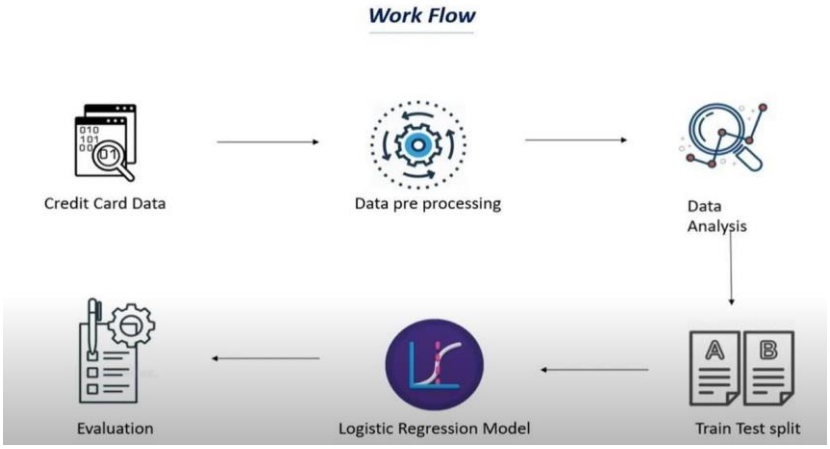

Architecture of the proposed credit card fraud detection framework.

Data Collection: The dataset utilized in this study was obtained from Kaggle and consists of credit card transaction records collected in September 2013 from European cardholders. As illustrated in Figure 3.2, the dataset comprises 284,807 transactions, among which 492 instances are annotated as fraudulent, reflecting a highly imbalanced class distribution typical of real-world fraud detection scenarios.

Preprocessing Step: Following data acquisition, preprocessing prepares transactions for analysis by addressing missing values, eliminating duplicates, and mitigating noise. Numerical attributes are normalized and scaled to comparable ranges to prevent dominance of any single feature, while categorical variables (e.g., merchant category, location) are encoded for model compatibility. These steps improve data consistency, stability, and reliability, providing a robust foundation for subsequent feature engineering and predictive modeling in fraud detection systems.

Feature Extraction: Feature extraction converts raw transactional data into meaningful attributes that inform fraud detection, including behavioral indicators, temporal patterns, transaction frequency, spending profiles, and geographic signals. Advanced systems incorporate time-series features to capture trends and anomalies over time. By emphasizing the most informative variables and reducing dimensionality, feature extraction enables learning algorithms to model relevant fraud patterns more effectively and improves predictive performance.

Data Prediction: Prediction constitutes the core of the fraud detection pipeline, wherein trained machine learning models classify transactions as legitimate or fraudulent using engineered features. Supervised classifiers such as Logistic Regression, Random Forest, and Support Vector Machines are employed to estimate the probability of fraud for each transaction. The model outputs calibrated risk scores, and transactions exceeding an operational threshold are flagged for intervention or manual review. This probabilistic decisioning enables risk-based actions while controlling false positives. In real-time deployments, predictions are generated with low latency to prevent fraudulent activity at the point of transaction, thereby reducing financial loss and improving customer security through timely detection and response.

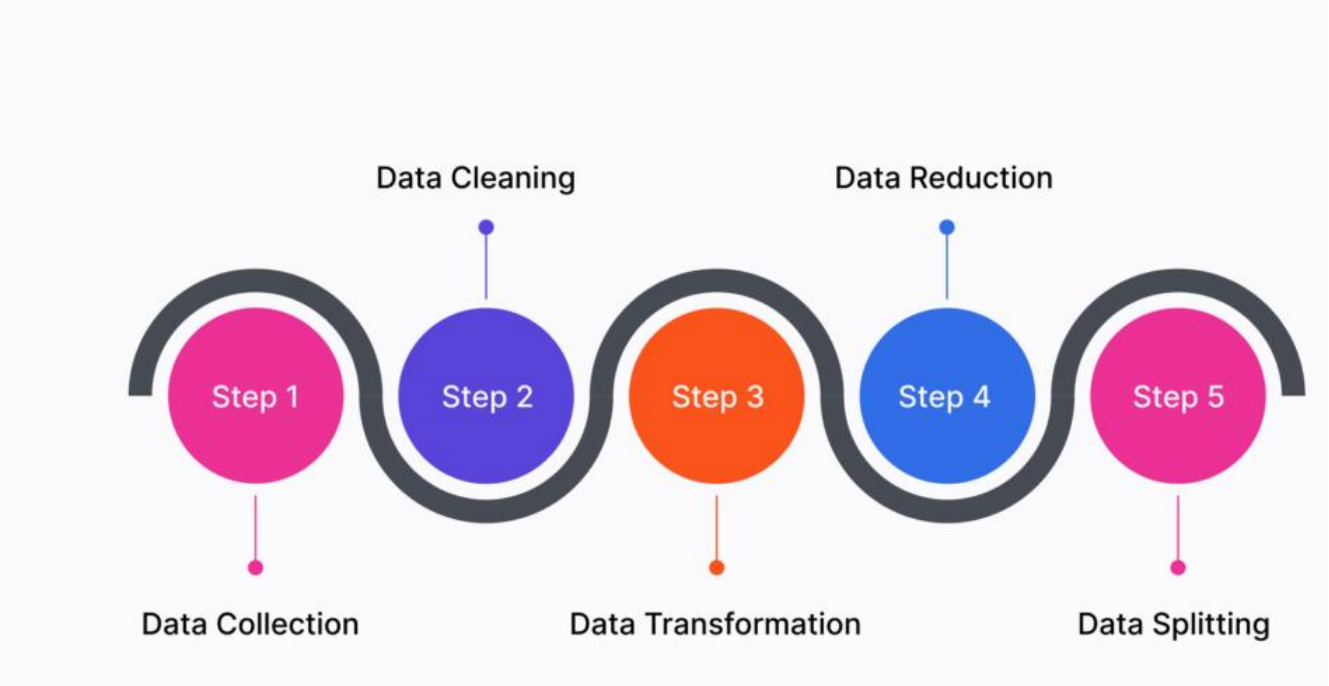

Data preprocessing workflow including collection, cleaning, transformation, reduction, and splitting.

Machine Learning Algorithms:

Logistic Regression: Logistic Regression is trained with class-weighted cross-entropy on standardized features and regularization to stabilize coefficients under multicollinearity. Calibrated probability outputs support risk scoring, with decision thresholds tuned to operational recall constraints. The model serves as a low-latency baseline for real-time screening and ensemble stacking.

Performance metrics of the Logistic Regression model (Class 1)

Random Forest: Random Forest is deployed using bootstrapped sampling and randomized feature subspaces to capture non-linear fraud signatures while controlling variance. Class-weighted training and

probability calibration are applied to mitigate skewed class priors. Aggregated tree probabilities provide stable risk scores and feature importance for governance and monitoring in real-time detection workflows.

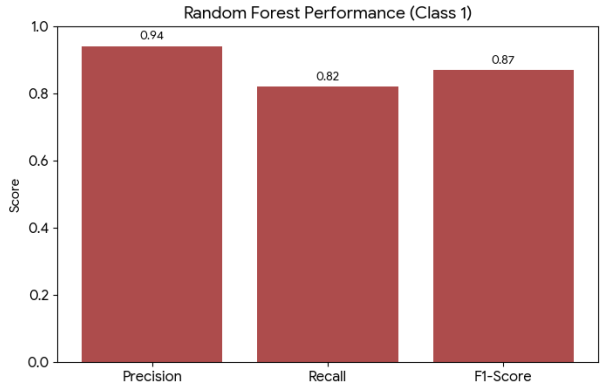

Performance metrics of the Random Forest model.

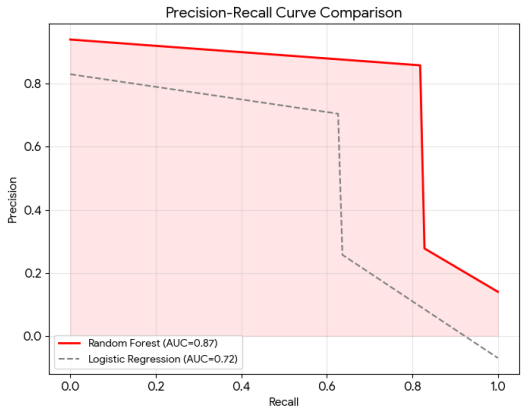

Precision–recall curve comparison of Logistic Regression and Random Forest models.

Xg-Boost: Xg-Boost is implemented with gradient-boosted trees optimized on imbalanced loss functions and regularization (L1/L2) to suppress overfitting on sparse fraud signals. Tree depth, learning rate, and scale_pos_weight are tuned to maximize PR-AUC. The model outputs calibrated fraud probabilities, enabling threshold-based real-time flagging under latency constraints.

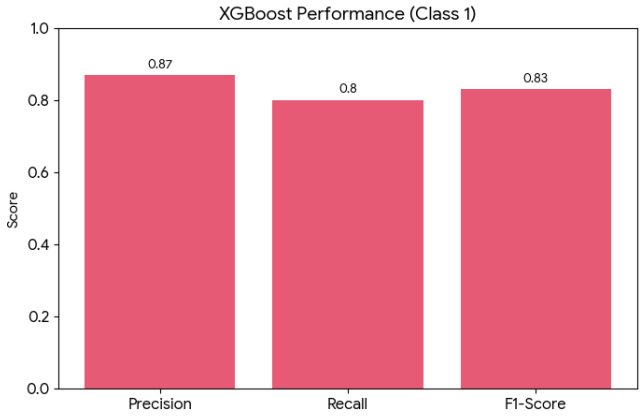

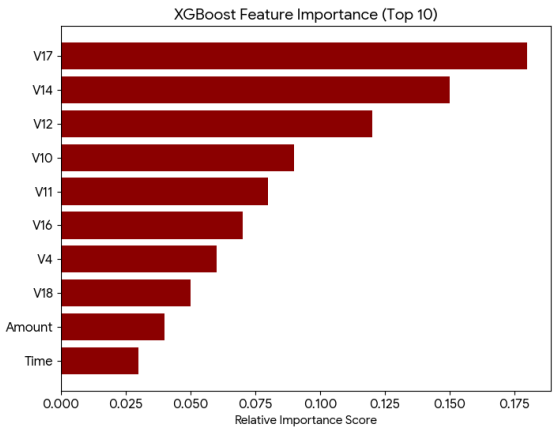

IV. RESULTS AND DISCUSSION

The comparative evaluation of Logistic Regression, Random Forest, and Xg-Boost highlights the impact of model capacity and class-imbalance handling on fraud detection performance. As shown in the class-wise metrics, all models achieve near-perfect overall accuracy due to the dominance of legitimate transactions; however, accuracy alone is not a reliable indicator in highly imbalanced settings. Therefore, the analysis focuses on precision, recall, and F1-score for the minority (fraud) class.

Logistic Regression demonstrates strong precision for fraudulent transactions (≈0.83) but relatively lower recall (≈0.63), indicating that while flagged transactions are often truly fraudulent, a substantial portion of fraud cases remains undetected. This behavior reflects the linear decision boundary and conservative probability thresholding, which favors minimizing false positives at the expense of missed fraud. Consequently, its F1-score (≈0.72) is lower than ensemble-based methods, limiting its suitability as a standalone detector in high-risk environments where recall is critical.

Random Forest achieves the most balanced performance among the evaluated models, with high precision (≈0.94) and recall (≈0.82) for the fraud class, resulting in the highest F1-score (≈0.87). The improvement can be attributed to variance reduction via bootstrapped aggregation and the ability to model non-linear interactions among transactional, temporal, and behavioral features. The ensemble effectively captures heterogeneous fraud signatures while maintaining stability under class imbalance through class-weighted training and probability aggregation. These properties make Random Forest particularly effective for operational fraud screening, where both detection coverage and false-alarm control are required.

Xg-Boost also demonstrates strong performance, achieving precision (≈0.87) and recall (≈0.80) for fraudulent transactions, with an F1-score (≈0.83). The boosted trees leverage additive modeling with regularization to learn complex fraud patterns from sparse minority signals. While slightly below Random Forest in F1-score, Xg-Boost provides competitive recall and can be tuned to prioritize detection sensitivity using imbalance-aware loss weighting (e.g., scale_pos_weight). Its calibrated probability outputs support threshold-based decisioning under real-time latency constraints.

```
----- Logistic Regression Results -----
Accuracy: 0.9991397773954567
              precision    recall  f1-score   support

           0       1.00      1.00      1.00     56864
           1       0.83      0.63      0.72        98

    accuracy                           1.00     56962
   macro avg       0.91      0.82      0.86     56962
weighted avg       1.00      1.00      1.00     56962
```

```
----- Random Forest Results -----
Accuracy: 0.9995962220427653
              precision    recall  f1-score   support

           0       1.00      1.00      1.00     56864
           1       0.94      0.82      0.87        98

    accuracy                           1.00     56962
   macro avg       0.97      0.91      0.94     56962
weighted avg       1.00      1.00      1.00     56962
```

```
----- XGBoost Results -----
Accuracy: 0.9994382219725431
              precision    recall  f1-score   support

           0       1.00      1.00      1.00     56864
           1       0.87      0.80      0.83        98

    accuracy                           1.00     56962
   macro avg       0.93      0.90      0.91     56962
weighted avg       1.00      1.00      1.00     56962
```

Classification report comparison of Logistic Regression, Random Forest, and XGBoost models.

Across all models, the classification reports indicate near-perfect weighted averages, which primarily reflect performance on the majority class and should not be over-interpreted. The minority-class metrics provide a more realistic assessment of fraud detection capability. The bar chart comparison further corroborates that ensemble methods outperform the linear baseline in recall and F1-score, demonstrating superior trade-offs between missed fraud and false alarms.

From a deployment perspective, Logistic Regression remains valuable as a low-latency baseline or as a component in stacked ensembles, whereas Random Forest offers the best operational balance for high-coverage fraud detection with interpretability via feature importance. XGBoost provides strong performance with fine-grained control over bias–variance trade-offs and is suitable for scenarios requiring aggressive recall optimization. Overall, the results validate the use of ensemble tree-based models with calibrated probabilities and threshold tuning as the most effective strategy for real-world credit card fraud detection under severe class imbalance.

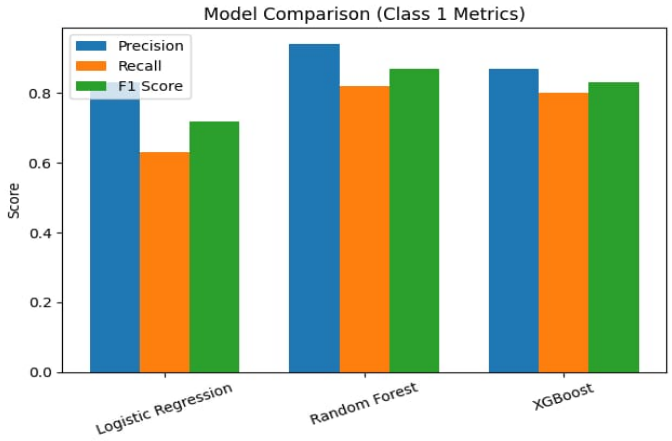

Comparison of precision, recall, and F1-score across models for Class 1

V. CONCLUSION AND FUTURE SCOPE

This study demonstrates the effectiveness of machine learning models for credit card fraud detection under extreme class imbalance. While Logistic Regression provides a fast and interpretable baseline, its lower recall indicates limited coverage of fraudulent activity. Ensemble methods significantly outperform the linear baseline by capturing non-linear fraud signatures and complex feature interactions. In particular, Random Forest achieves the most balanced trade-off between precision and recall, resulting in superior F1-score and operational reliability. Xg-Boost delivers competitive performance with strong recall and tunable bias–variance characteristics, making it suitable for scenarios requiring aggressive detection sensitivity. The findings confirm that ensemble tree-based models with calibrated probability outputs and threshold optimization are well-suited for real-time fraud screening pipelines, where minimizing financial loss while controlling false positives is critical.

Based on these results, practical improvements to the deployed fraud detection system include adopting an ensemble-first strategy for primary screening, complemented by Logistic Regression as a low-latency fallback or stacking component. Thresholds should be optimized using cost-sensitive objectives aligned with business risk (e.g., higher penalty for false negatives). Continuous monitoring for concept drift is essential due to evolving fraud tactics; therefore, periodic retraining and champion–challenger evaluation should be integrated into the production workflow. Probability calibration and PR-AUC–centric evaluation should be maintained to ensure stable decisioning under class imbalance. Incorporating feature attribution (e.g., tree-based importance) can further enhance analyst trust and governance.

For future and related projects, extending the system to support streaming ingestion and real-time feature computation will improve deployment readiness. Investigating hybrid ensembles (e.g., stacking Random Forest and Xg-Boost) and semi-supervised or anomaly detection methods can improve early detection of novel fraud patterns. Drift-aware learning, online retraining, and adaptive thresholds can further sustain performance over time. Beyond card fraud, the same framework can be transferred to adjacent domains such as digital payment abuse, insurance claim fraud, transaction laundering detection, and account takeover prevention by adapting feature engineering and cost functions to domain-specific risk profiles.

The proposed framework delivers high detection performance with governance-friendly interpretability and scalability. Future work includes concept-drift monitoring and online learning for adaptive defense.

REFERENCES

1. **Bhattacharyya, S.** (2011). Credit Card Fraud Detection Using Data Mining Techniques. CSSE Proceedings.
2. **Chandran, P. R.** Credit Card Fraud Detection System Using Machine Learning Algorithms: A Survey. JEST.
3. **Han, J., & Kamber, M.** (2012). Data Mining: Concepts and Techniques (3rd ed.). Elsevier.
4. **Kotsiantis, S. B.** (2006). A review of Classification and Combining Techniques. Artificial Intelligence Review.
5. **Monika, R., & Shukla, M.** (2019). Credit Card Fraud Detection using Machine Learning Techniques. ICACDS.
6. **Shahrukh, M.** Application of Support Vector Machines for Credit Card Fraud Detection. Journal of Computer Science.
7. **Waleed, N.** Exploring the Use of Logistic Regression for Fraud Detection in Financial Transactions. IJACSA.
8. **He, H., & Garcia, E. A.** (2009). Learning from Imbalanced Data. IEEE Transactions on Knowledge and Data Engineering.

9. **Chen, T., & Guestrin, C.** (2016). XGBoost: A Scalable Tree Boosting System. ACM SIGKDD International Conference.
10. **Breiman, L.** (2001). Random Forests. Machine Learning Journal.
11. **Fawcett, T.** (2006). An introduction to ROC analysis. Pattern Recognition Letters.
12. **Quah, J. T., & Sriganesh, M.** (2008). Real-time credit card fraud detection using computational intelligence. Expert Systems with Applications.
13. **Dal Pozzolo, A., et al.** (2014). Learned lessons in credit card fraud detection from a practitioner perspective. Expert Systems with Applications.