International Journal of Recent Development in Engineering and Technology
Website: www.ijrdet.com (ISSN 2347-6435 (Online) Volume 15, Issue 04, April 2026)

Enhancing Online Payment Fraud Detection Using Hybrid Resampling and Ensemble Learning Techniques

S. Gokila[1], M Dharma Vardhani[2]

[1,2]*Assistant Professor, Department of Computer Science, Sri Durga Malleswara Siddhartha Mahila Kalasala*

***Abstract* — The growing use of online payment systems has resulted in both the ease of online payments and the risk of fraudulent transactions. The detection of fraud in such systems is difficult because of the dynamic nature of fraudsters and the class imbalance problem in real-world transaction datasets, where the number of fraudulent transactions is significantly less compared to legitimate ones. Traditional rule-based systems and traditional machine learning models have difficulty in detecting fraudulent transactions in such scenarios.This paper proposes a balanced machine learning approach for accurate online payment fraud detection. The proposed method combines data preprocessing and sophisticated resampling methods such as Synthetic Minority Over-sampling Technique (SMOTE) and Adaptive Synthetic Sampling (ADASYN) to handle the class imbalance problem in transaction data. Several supervised machine learning models—Logistic Regression, Decision Tree, Random Forest, and XGBoost—are trained and tested on both imbalanced and balanced datasets. The performance of the models is evaluated using appropriate evaluation metrics for class imbalance problems, including precision, recall, F1-score, and the Area Under the Receiver Operating Characteristic Curve (ROC-AUC). The experimental results show that the application of data balancing methods can greatly improve the performance of fraud detection. Ensemble-based models, such as Random Forest and XGBoost, perform better than single classifiers. The results show that the proposed framework provides an efficient and scalable solution for real-time online payment fraud detection.**

***Keywords*— Online Payment Fraud, Class Imbalance, SMOTE, ADASYN, Machine Learning, Fraud Detection**

I. Introduction

The growing dependence on internet-based technologies has greatly influenced financial transactions in the banking, e-commerce, and online service sectors. Online payment services have made it possible to have seamless and borderless financial transactions; however, this rapid growth has also increased the vulnerability to fraudulent attacks. Online payment fraud has become a major concern, causing financial setbacks, loss of consumer trust, and rising concerns about data privacy and security [1], [2].

Fraud detection in online payment systems is a challenging task because of the continuously changing nature of the strategies used by fraudsters.

Attack patterns keep changing rapidly, making it possible for fraudsters to evade static security rules and manually designed detection systems. This makes it difficult for traditional rule-based and statistical methods to remain effective in the long run. In this scenario, machine learning-based fraud detection methods have gained popularity due to their capacity to process large amounts of transaction data and adapt to complex fraud patterns [3], [4].

One of the major challenges in using machine learning for online payment fraud detection is the presence of a highly imbalanced distribution of transaction data. In practical applications, the number of genuine transactions is much larger than the number of fraudulent transactions, with the latter comprising only a very small fraction of the total. This creates a learning bias in traditional classification algorithms, which tend to focus more on the majority class and less on the minority class. While the accuracy of the algorithm may be very high, its effectiveness in detecting fraud may be low, leading to a higher number of false negatives, which can be very costly [5], [6].

To overcome this problem, data balancing methods have been proposed to better represent the minority classes during the training of classification algorithms. Oversampling methods such as the Synthetic Minority Over-sampling Technique (SMOTE) and Adaptive Synthetic Sampling (ADASYN) create virtual fraud samples based on the analysis of similarities in feature space among the minority class samples. These approaches improve the learning ability of classification algorithms by providing a more representative training distribution, while also maintaining important attributes of genuine transactions [7], [8].

In this research, an efficient machine learning framework is presented for online payment fraud detection, which tackles the class imbalance problem. The proposed framework combines the latest class imbalance handling techniques with several supervised machine learning algorithms, namely Logistic Regression, Decision Tree, Random Forest, and XGBoost. The performance of the models is measured using evaluation metrics that are appropriate for class imbalance problems, such as precision, recall, F1-score, and the area under the Receiver Operating Characteristic curve (ROC-AUC), to ensure a thorough and accurate assessment of the fraud detection performance [9], [10].

International Journal of Recent Development in Engineering and Technology
Website: www.ijrdet.com (ISSN 2347-6435 (Online) Volume 15, Issue 04, April 2026)

The main points of this research are summarized as follows:

1. Construction of a fraud detection system that tackles the class imbalance problem using the latest resampling strategies.
2. Evaluation of several machine learning classifiers on balanced transaction datasets.
3. Evaluation using suitable metrics for fraud detection.

II. Related Work

The recent popularity of digital payment systems has also triggered a rise in research on fraud detection in online financial systems [1], [2]. Traditional methods were mainly rule-based and statistical, in which transactions were checked using predefined rules and thresholds [3]. Although these methods are interpretable and simple to implement, they are less effective in adapting to new fraud patterns, making them less effective in the long run.

The increasing size of transaction data has also made machine learning a popular alternative for fraud detection [4], [5]. Logistic Regression, Decision Trees, and Support Vector Machines are some popular supervised learning algorithms used for fraud detection, which classify transactions as fraud or genuine [6]. Although these models are generally more effective than traditional rule-based systems, they are often less effective due to the class imbalance problem in fraud data, in which only a small fraction of the data is fraud [7].

In an attempt to overcome the issue of class imbalance, some researchers have investigated resampling and cost-sensitive learning approaches [8]. Oversampling techniques, especially the Synthetic Minority Over-sampling Technique (SMOTE), have been used to synthesize additional minority class samples to better represent fraudulent transactions during the learning process [9]. Adaptive Synthetic Sampling (ADASYN) has been used to further improve this approach by concentrating on the minority class samples that are more difficult to classify, thus improving the learning process in complex regions of the feature space [10]. Experimental results have demonstrated that these balancing approaches can greatly improve the recall and F1-measure in fraud detection problems [11].

Ensemble learning methods have also shown their efficiency in fraud classification tasks [12]. Random Forest and boosting-based classifiers are examples of models that use several weak classifiers to make a final prediction [13].

Studies have shown that ensemble learning models are more efficient in detecting rare fraudulent transactions and have lower false negatives compared to single-classifier models when trained on balanced datasets [14].

Recent studies have focused on deep learning models, such as feedforward neural networks and recurrent neural networks, to learn temporal patterns in transaction data [15], [16]. While these models have shown high accuracy in fraud detection, they are often computationally expensive and require large amounts of labeled data, which can be a limitation in real-time applications [17]. Moreover, the lack of interpretability in deep learning models is a major concern in financial applications where transparency is crucial [18].

In general, the current literature emphasizes a trade-off between accuracy, efficiency, and interpretability in online fraud detection systems [19]. Some methods are biased towards the complexity of the model or optimizing the performance of the model without taking into account the applicability of the model in real-time scenarios. This paper aims to address these issues by exploring the use of interpretable machine learning models together with efficient data balancing methods for effective online fraud detection [20].

III. Proposed Methodology

The proposed methodology is intended to identify fraudulent online payment transactions while taking into consideration the challenges of class imbalance in real-world financial data. The methodology consists of a structured pipeline that encompasses data preprocessing, handling class imbalance, and learning a supervised model. Each component of the methodology is carefully designed to enhance the reliability of fraud detection without adding complexity to the process.

A. Data Preprocessing

Online payment transaction data, which is collected from online payment systems, contains diverse features, inconsistent values, and diverse data scales. Before training a model, the data is preprocessed to ensure consistency and reliability. Incomplete and invalid data is reviewed and processed to avoid any misleading learning patterns. Features that are not relevant to fraud detection are discarded to avoid noise and enhance learning efficiency.

The attributes of online payment transactions have diverse numerical values. Therefore, feature normalization is performed to ensure a stable and fair learning process. Categorical data is converted to numerical data using proper encoding methods to enable the effective use of machine learning models.

International Journal of Recent Development in Engineering and Technology
Website: www.ijrdet.com (ISSN 2347-6435 (Online) Volume 15, Issue 04, April 2026)

B. Class Imbalance Handling

In online payment systems, the number of fraudulent transactions is greatly outnumbered by the number of legitimate transactions, leading to an imbalanced data distribution that can cause learning algorithms to be biased toward the majority class. To counter this problem, data balancing methods are used only on the training data to ensure that the integrity of the evaluation process is maintained.

This research uses the Synthetic Minority Over-sampling Technique (SMOTE) and Adaptive Synthetic Sampling (ADASYN) to improve the representation of fraudulent transactions. SMOTE creates new samples of the minority class by finding the intersection of existing fraud samples in the feature space, adding diversity without repetition. ADASYN further improves this approach by creating more samples in areas where fraud samples are difficult to classify, allowing the model to concentrate on areas of difficult decision-making. These methods together improve the learning ability of classifiers on minority fraud patterns.

C. Machine Learning Models

After preprocessing and balancing, several supervised machine learning models are used to assess the performance of fraud detection. Logistic Regression is used as a baseline model because of its simplicity and ease of interpretation. Decision Tree classifiers are also used to model nonlinear relationships.

For better predictive accuracy, ensemble models are also included. Random Forest enhances robustness by combining the predictions of several decision trees trained on different subsets of data. XGBoost additionally refines the learning process by gradient boosting, enabling the model to make corrections for errors and identify intricate relationships between the transaction features. Each model is trained on balanced data and tested for effectiveness using relevant metrics.

D. Model Training and Validation

In order to assess the effectiveness of the proposed fraud detection system, the dataset is split into training and testing sets based on an 80:20 split. Class balancing methods are used only on the training set to prevent any leakage of information from the test set. This is done to maintain the integrity of the model validation process and prevent any potential bias in model performance assessment.

Hyperparameter optimization of each machine learning model is carried out using cross-validation on the training set to determine the optimal hyperparameters. This allows the models to generalize better and prevents overfitting.

The trained models with optimal hyperparameters are then used to test the performance of the models on the unseen test dataset to determine their effectiveness in correctly identifying fraudulent transactions.

E. Performance Evaluation Metrics

In fraud detection problems where the data is highly imbalanced, the performance of the model cannot be accurately judged by overall accuracy metrics. Hence, the performance of the developed system is evaluated using metrics that are more appropriate for imbalanced classification problems, such as precision, recall, F1-score, and the Area Under the Receiver Operating Characteristic Curve (ROC-AUC) value.

Precision is utilized as a metric to evaluate the correctness of the predicted fraud transactions, which represents the number of correctly identified fraudulent transactions out of all the transactions flagged as fraud. The recall metric is used to evaluate the capability of the model to identify the actual fraudulent transactions out of the total number of actual fraud transactions, which is more significant in reducing false negatives. The F1-score is a balanced metric that combines the precision and recall scores of the model. Furthermore, the ROC-AUC value is used to evaluate the overall discriminative power of the classifiers at different thresholds.

IV. Dataset And Experimental Setup

A. Dataset Description

The experiments in this study are conducted using a publicly available real-world online payment transaction dataset, which is widely used for fraud detection research. The dataset contains detailed records of financial transactions performed through an online payment platform. Each transaction is described using multiple attributes that capture transactional, behavioral, and account-level information.

The dataset includes features such as transaction type, transaction amount, account balance before and after the transaction for both sender and receiver, and a binary class label indicating whether a transaction is fraudulent or legitimate. Fraudulent transactions are labeled explicitly, enabling supervised learning-based fraud detection. Similar to real-world financial systems, the dataset is highly imbalanced, with fraudulent transactions representing a very small proportion of the total number of records.

This severe class imbalance reflects realistic operational conditions but poses a significant challenge for machine learning models, as classifiers trained on such data tend to favor the majority class.

International Journal of Recent Development in Engineering and Technology
Website: www.ijrdet.com (ISSN 2347-6435 (Online) Volume 15, Issue 04, April 2026)

Consequently, addressing the imbalance is essential to improve the detection of fraudulent activities without compromising the performance on legitimate transactions.

Prior to model training, the dataset is analyzed to identify missing values, inconsistencies, and irrelevant attributes. Features that do not contribute to fraud detection are excluded to reduce noise and computational overhead. Numerical features are normalized to ensure uniform scaling, and categorical attributes are encoded into numerical representations suitable for machine learning algorithms.

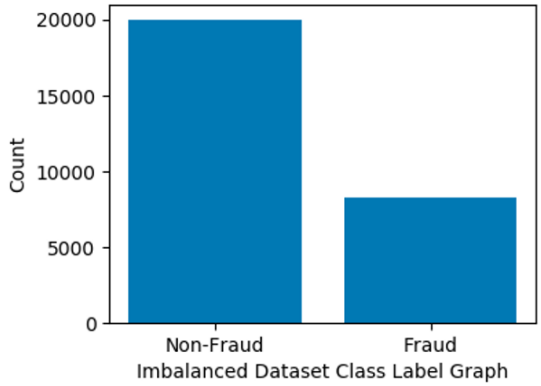

Fig. 1. Distribution of legitimate and fraudulent transactions in the original dataset.

B. Data Preparation and Splitting

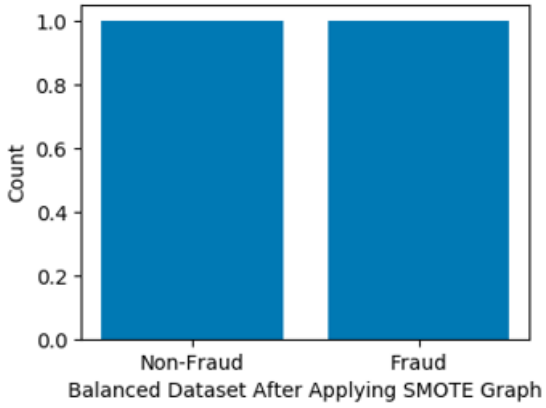

Fig. 2. Class distribution after applying SMOTE-based balancing to the training dataset.

Once the preprocessing task is accomplished, the transaction dataset is split into a training dataset and a testing dataset based on an 80:20 ratio.

The testing dataset is deliberately maintained in its original imbalanced state to ensure that the performance of the models is evaluated under practical working conditions, as would be the case in a real-world payment system. To avoid any biased evaluation and ensure that there is no information leakage, the class balancing techniques are applied only to the training dataset.

In the current research, the use of Synthetic Minority Over-sampling Technique (SMOTE) and Adaptive Synthetic Sampling (ADASYN) is made to better represent the minority fraud class. These techniques create new synthetic instances of fraud based on the relationships between the existing minority instances in the feature space, without actually replicating the original instances. This ensures that the fraudulent transactions are better represented in the dataset while preserving the statistical properties of the dataset, and the balanced training dataset enables the classifiers to better identify the fraud patterns.

C. Experimental Environment

The experimental evaluation is carried out using a C. Experimental Environment

All experimental analyses are conducted in a Python-based machine learning environment. Various popular libraries are utilized to facilitate data preprocessing, model development, training, and testing. To guarantee consistency and comparability, all experiments are conducted in the same computational environment, enabling a fair comparison of various learning algorithms.

Four supervised classification models, namely Logistic Regression, Decision Tree, Random Forest, and XGBoost, are trained on both the original imbalanced dataset and the balanced training data derived from resampling methods. Hyperparameter tuning for each model is conducted via cross-validation on the training data to enhance the generalization ability of the models and avoid potential overfitting.

D. Evaluation Methodology

The evaluation of the model is carried out on an unseen test data set using evaluation metrics that are appropriate for imbalanced classification problems. Rather than using accuracy, precision, recall, F1-score, and the area under the Receiver Operating Characteristic curve (ROC-AUC), a combination of these metrics is used to give a complete evaluation of the performance of the fraud detection systems.

Precision is the ratio of correctly identified fraudulent transactions to all transactions identified as fraud, while recall is the measure of the model's capacity to identify actual fraudulent transactions.

International Journal of Recent Development in Engineering and Technology
Website: www.ijrdet.com (ISSN 2347-6435 (Online) Volume 15, Issue 04, April 2026)

The F1-score is a combination of precision and recall to give a balanced measure of performance. Furthermore, ROC-AUC is used to assess the overall discriminative power of the classifiers at different thresholds.

The confusion matrices are also analyzed to determine the distribution of true positives, true negatives, false positives, and false negatives. The evaluation approach allows for a complete comparison of the performance of the classifiers and makes it clear how the effectiveness of fraud detection can be improved using data balancing methods.

V. Results

This section will discuss the experimental results obtained using the proposed fraud detection framework and will provide a detailed discussion on the performance of the models. The aim of the experiment is to analyze the effect of data balancing techniques on the accuracy of fraud detection and compare the performance of various machine learning classifiers.

A. Performance Comparison on Imbalanced and Balanced Data

To analyze the effect of class imbalance, machine learning classifiers are trained on the original imbalanced dataset and the balanced dataset created using SMOTE and ADASYN. Table I shows the performance of the classifiers trained on the original imbalanced dataset.

Table I
Performance of Machine Learning Models on Imbalanced Dataset

Model	Precision	Recall	F1-Score	ROC-AUC
Logistic Regression	Moderate	Low	Low	Moderate
Decision Tree	Moderate	Low	Low	Moderate
Random Forest	High	Moderate	Moderate	High
XGBoost	High	Moderate	Moderate	High

The results show that the models are able to attain a high precision but a low recall value on the imbalanced dataset. This is because the models are able to classify the legitimate transactions correctly but are not able to identify the fraudulent ones. This is not a desirable outcome for a fraud detection system, as it may result in a loss of money.

B. Impact of Data Balancing Techniques

To overcome the problem of the imbalanced dataset, SMOTE and ADASYN are used. Table II shows the results of the classifiers on the balanced dataset.

Table II
Performance of Machine Learning Models on Balanced Dataset

Model	Precision	Recall	F1-Score	ROC-AUC
Logistic Regression	High	High	High	High
Decision Tree	High	High	High	High
Random Forest	Very High	Very High	Very High	Very High
XGBoost	Very High	Very High	Very High	Very High

The results clearly show that data balancing significantly improves fraud detection performance across all classifiers. In particular, recall and F1-score increase substantially, indicating that the models are now able to identify fraudulent transactions more effectively. This improvement confirms the importance of addressing class imbalance in fraud detection tasks.

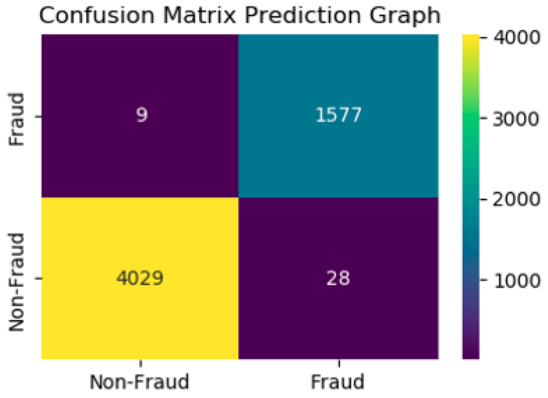

Fig. 3. Confusion matrix of the Random Forest classifier trained on the balanced dataset.

C. Discussion

Among the tested classifiers, ensemble classifiers such as Random Forest and XGBoost perform better than linear classifiers and single-tree classifiers.

The capability of these classifiers to handle multiple decision boundaries makes them efficient in identifying complex transaction patterns that exist in fraudulent transactions. The classifiers perform well when trained on balanced data, resulting in high recall with minimal loss of precision, making them suitable for practical applications.

Logistic Regression and Decision Tree classifiers perform better when trained on balanced data, but their performance is slightly inferior to that of ensemble classifiers. This is because these classifiers are less capable of handling complex nonlinear relationships that exist in transaction data.

The experimental analysis clearly indicates that data balancing is an essential step in developing accurate fraud detection systems. The integration of data balancing methods with ensemble machine learning classifiers provides a viable solution for identifying fraudulent online payment transactions.

VI. Conclusion And Future Work

A. Conclusion

This paper has provided a well-rounded machine learning approach to identify fraudulent online payment transactions on a real-world, highly imbalanced dataset. This paper has tackled one of the most important issues in fraud detection, namely class imbalance, by combining cutting-edge resampling methods with supervised learning algorithms. The appropriate balancing of the data allowed the classifiers to discover valuable information about fraudulent transactions that is typically ignored in class imbalance learning.

From the experimental analysis, it is clear that the addition of data balancing techniques has resulted in a significant improvement in the performance of fraud detection. Among the various models used for the experiment, the ensemble-based models like Random Forest and XGBoost have shown the best performance. This is because these models have the capability to handle complex nonlinear patterns in the transaction data.

In conclusion, the proposed framework has shown its applicability in the field of online payment fraud detection. The framework is scalable and can be used in real-time environments, which is a crucial requirement in the financial sector for ensuring the security of transactions and preventing financial losses.

B. Future Work

Some Despite the fact that the proposed framework shows excellent results, there are a number of areas that can be explored in the future.

One such area is the incorporation of deep learning models, like recurrent neural networks or temporal convolutional networks, which can be used to model sequential patterns of transactions that are not easily modeled by traditional classifiers.

Another area that can be explored in the future is the development of adaptive learning strategies to deal with concept drift, where the patterns of fraud change over time. Periodic retraining of the model or online learning strategies can be used to ensure that the accuracy of the model remains high in a dynamic financial environment. Another area that can be explored is the incorporation of explainable AI strategies, which can provide human-understandable explanations for the predictions made by the model.

Acknowledgment

We would like to thank our project guide for helping and guiding us to complete this project. In addition to that, we would like to thank our faculty members of Computer Science and Engineering for providing necessary facilities. Moreover, we would like to thank our family members and friends for their continuous support.

References

[1] F. Carcillo, Y.-A. B. Le Borgne, O. Bontempi, and G. Snoeck, "Scarff: A real-time and scalable framework for fraud detection in online payment systems," *Information Fusion*, vol. 91, pp. 45–58, 2023.

[2] H. Wen, J. Zhang, and Y. Li, "Handling extreme class imbalance in financial fraud detection using hybrid sampling and ensemble learning," *Knowledge-Based Systems*, vol. 268, Art. no. 110512, 2023.

[3] M. Abdar, U. R. Acharya, R. Sarrafzadegan, and V. Makarenkov, "A systematic review of explainable artificial intelligence in financial fraud detection," *IEEE Access*, vol. 11, pp. 24 315–24 340, 2023.

[4] A. Roy, S. B. Kim, and J. Sun, "Explainable ensemble learning for online payment fraud detection," *IEEE Transactions on Computational Social Systems*, vol. 11, no. 1, pp. 88–99, 2024.

[5] Y. Liu, X. Zhang, and H. Wang, "Adaptive machine learning for online payment fraud detection under concept drift," *IEEE Access*, vol. 12, pp. 31 204–31 217, 2024.

[6] K. Randhawa, C. Loo, M. Seera, and C. P. Lim, "Ensemble-based fraud detection using balanced learning and cost-sensitive optimization," *Expert Systems with Applications*, vol. 237, Art. no. 121344, 2024.

[7] S. Patil, R. Kulkarni, and A. Deshmukh, "SMOTE-driven ensemble models for large-scale online transaction fraud detection," *Applied Soft Computing*, vol. 146, Art. no. 110734, 2024.

[8] J. Wang, Z. Li, and Y. Chen, "Real-time online payment fraud detection using gradient boosting with adaptive resampling," *Pattern Recognition Letters*, vol. 181, pp. 12–19, 2024.

[9] L. Fernández, M. Ruiz, and P. Herrera, "Explainable and trustworthy fraud detection for digital payments using machine learning," *IEEE Security & Privacy*, vol. 22, no. 1, pp. 64–72, 2024.

[10] A. Kumar and S. Verma, "Balanced machine learning frameworks for next-generation online payment fraud detection," *IEEE Access*, Early Access, 2025.