International Research Journal of Innovations in Engineering and Technology (IRJIET)

ISSN (online): 2581-3048

Volume 10, Issue 4, pp 61-65, April-2026

https://doi.org/10.47001/IRJIET/2026.104007

Credit Card Fraud Detection: Mitigating Extreme Class Imbalance Using Synthetic Oversampling and Ensemble Machine Learning

[1]Prathmesh Sunil Dhobe, [2]Anjaneya Kokre, [3]Suyash Kale, [4]Kartik Sabe, [5]Shahrukh Shaikh

[1,2,3,4]Student, Department of Artificial Intelligence & Machine Learning, Ajeenkya D.Y. Patil School of Engineering, Maharashtra, India

[5]Guide / Supervisor, Professor, Department of Artificial Intelligence & Machine Learning, Ajeenkya D.Y. Patil School of Engineering, Maharashtra, India

Abstract **- The rapid proliferation of digital payment infrastructure has established credit card transactions as the backbone of the modern global economy, concurrently exposing financial networks to sophisticated fraudulent activities. The automated detection of such anomalies presents a significant algorithmic challenge due to extreme class imbalance, as fraudulent instances typically represent less than 0.5% of the overall transaction volume. This research proposes a robust, machine learning-based classification architecture utilizing a highly imbalanced dataset of 284,807 transactions, where the minority fraud class constitutes merely 0.17% of the data. To neutralize the statistical bias introduced by this skew, rigorous data preprocessing techniques including Z-score standardization and stratified splitting were implemented. The Synthetic Minority Over-sampling Technique (SMOTE) was deployed strictly within the training environment to synthetically balance the class distributions and prevent algorithmic convergence toward majority-class predictions. A comparative analysis was conducted evaluating a linear Logistic Regression classifier against a non-linear Random Forest ensemble. Empirical analysis demonstrates that while the linear model achieved high theoretical class separation, the Random Forest ensemble delivered superior operational performance. By optimizing the precision-recall trade-off, achieving a precision of 0.84 and a recall of 0.83, the ensemble model successfully minimized false negative rates without inflating false positive rates, proving its viability for real-world deployment in institutional financial security systems.**

Keywords: Credit Card Fraud, Class Imbalance, SMOTE, Random Forest, Logistic Regression, Anomaly Detection.

I. INTRODUCTION

The transition toward cashless societies has resulted in an exponential increase in digital transaction volumes. While this paradigm shift offers unprecedented efficiency, it simultaneously exposes financial institutions and consumers to severe cyber-financial threats. Credit card fraud inflicts substantial financial damage on banking institutions worldwide, making transaction-level interception critical to minimizing direct capital attrition and maintaining consumer trust. Historically, financial institutions relied on deterministic, rule-based expert systems to flag anomalous transactions based on static geographical or velocity parameters. However, malicious actors continuously alter their methodologies to bypass these static security protocols, rendering heuristic engines obsolete and highly prone to generating unsustainable false-positive rates. Machine learning models offer a dynamic alternative, capable of autonomously identifying hidden, non-linear relationships within complex transaction data.

The central challenge in deploying predictive modeling for financial fraud detection is the phenomenon of extreme class imbalance. In real-world datasets, fraudulent instances are microscopically rare compared to legitimate consumer behavior. Standard machine learning classifiers, which optimize for global accuracy, fall victim to the accuracy paradox. Such models invariably predict the majority class in all instances, achieving over 99% accuracy while failing entirely at isolating actual fraudulent behavior. Previous literature in anomaly detection frequently attempted to resolve this through random undersampling, which discards critical behavioral variance, or naive oversampling, which exacerbates algorithmic overfitting.

This paper formulates a mathematical and algorithmic solution to neutralize predictive bias. By integrating the Synthetic Minority Over-sampling Technique (SMOTE) with robust machine learning architectures, the proposed framework maximizes fraud detection sensitivity while maintaining strict operational exactness. The primary objective is to evaluate the efficacy of linear and ensemble approaches on synthetically balanced data to determine the

International Research Journal of Innovations in Engineering and Technology (IRJIET)

ISSN (online): 2581-3048

Volume 10, Issue 4, pp 61-65, April-2026

https://doi.org/10.47001/IRJIET/2026.104007

optimal deployment architecture for minimizing both financial loss and customer friction.

1.1 Project Overview

The transition toward cashless societies has resulted in an exponential increase in digital transaction volumes. While this paradigm shift offers unprecedented efficiency, it simultaneously exposes financial institutions and consumers to severe cyber-financial threats. Credit card fraud inflicts substantial financial damage on banking institutions worldwide, making transaction-level interception critical to minimizing direct capital attrition and maintaining consumer trust. Historically, financial institutions relied on deterministic, rule-based expert systems to flag anomalous transactions based on static geographical or velocity parameters. However, malicious actors continuously alter their methodologies to bypass these static security protocols, rendering heuristic engines obsolete and highly prone to generating unsustainable false-positive rates. Machine learning models offer a dynamic alternative, capable of autonomously identifying hidden, non-linear relationships within complex transaction data.

1.2 Problem Statement & Need of the Project

The central challenge in deploying predictive modeling for financial fraud detection is the phenomenon of extreme class imbalance. In real-world datasets, fraudulent instances are microscopically rare compared to legitimate consumer behavior. Standard machine learning classifiers, which optimize for global accuracy, fall victim to the accuracy paradox. Such models invariably predict the majority class in all instances, achieving over 99% accuracy while failing entirely at isolating actual fraudulent behavior.

This paper formulates a mathematical and algorithmic solution to neutralize predictive bias. By integrating the Synthetic Minority Over-sampling Technique (SMOTE) with robust machine learning architectures, the proposed framework maximizes fraud detection sensitivity while maintaining strict operational exactness. The primary objective is to evaluate the efficacy of linear and ensemble approaches on synthetically balanced data to determine the optimal deployment architecture for minimizing both financial loss and customer friction.

II. LITERATURE REVIEW AND METHODOLOGY

This section reviews the historical and algorithmic context of financial anomaly detection, followed by the specific methodological architecture implemented in this study to resolve the identified challenges of class imbalance and data leakage.

2.1 Extant Literature and Methodological Challenges

The pursuit of automated fraud detection has driven significant academic research, evolving from simple statistical thresholds to advanced artificial intelligence paradigms. Early implementations relied on velocity checks, which often produced unsustainable false-positive rates [8]. In predictive modeling, learning from datasets where the target class is severely skewed represents a fundamental challenge. Traditional random undersampling discards critical behavioral variance, while naive random oversampling exacerbates algorithmic overfitting [7]. A review of recent implementations reveals recurring methodological flaws, most notably "data leakage," which occurs when oversampling techniques are applied prior to train-test splitting [4]. This study addresses these critical gaps by strictly isolating the test set and prioritizing operational metrics over misleading global accuracy scores [6].

2.2 Dataset and Exploratory Analysis

The empirical foundation of this study is a publicly available dataset comprising 284,807 European credit card transactions. Due to strict financial confidentiality regulations, the original variables were transformed via Principal Component Analysis (PCA), yielding 28 numerical features. Initial statistical profiling confirmed the extreme class imbalance, with only 492 transactions (0.17%) officially labeled as fraudulent. Exploratory data analysis revealed distinct statistical outliers within the monetary values, as fraudulent transactions possess a higher average fiat value compared to legitimate transactions. Furthermore, Kernel Density Estimate distributions demonstrated stark divergence in the probability distributions between classes for specific principal components, indicating high predictive power.

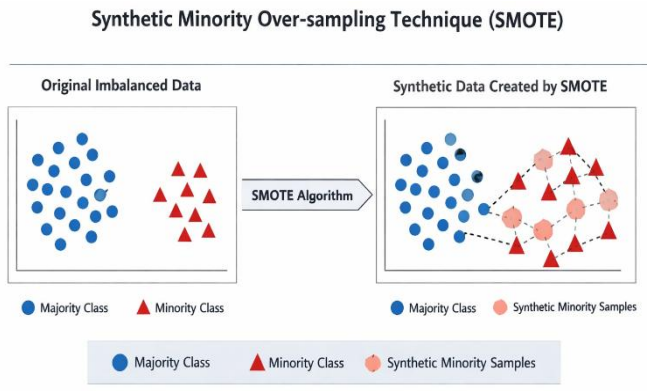

Figure 1: Visualizing Data Balancing using the Synthetic Minority Over-sampling Technique (SMOTE)

International Research Journal of Innovations in Engineering and Technology (IRJIET)

ISSN (online): 2581-3048

Volume 10, Issue 4, pp 61-65, April-2026

https://doi.org/10.47001/IRJIET/2026.104007

2.3 Algorithmic Balancing and Predictive Modeling Architecture

To prevent non-transformed variables from mathematically dominating the PCA features, Z-score standardization was applied. To strictly prevent data leakage, the dataset was first split using a stratified 80/20 ratio. To rectify the extreme class imbalance, the Synthetic Minority Over-sampling Technique (SMOTE) [2] was deployed exclusively on the training subset, balancing the data mathematically without exposing the model to the unseen testing data. The modeling phase then evaluated two distinct algorithms. Logistic Regression was deployed to establish a probabilistic linear baseline. To capture more complex, non-linear fraud vectors, the Random Forest Classifier was implemented [3]. This ensemble method utilizes bootstrap aggregating to build multiple decision trees on random subsets of data, relying on a majority vote to minimize individual tree variance and prevent overfitting.

III. RESULTS AND DISCUSSIONS

Given the severe class imbalance inherent in financial datasets, global accuracy provides a highly misleading representation of operational efficacy. The evaluation framework therefore prioritized specific metrics derived from the confusion matrix. Precision was utilized to measure the system's exactness, indicating the proportion of positive identifications that were actually correct. Recall, or sensitivity, measured the system's completeness by calculating the proportion of actual fraudulent transactions successfully identified. The F1-score provided the harmonic mean of precision and recall, while the Area Under the Receiver Operating Characteristic Curve (ROC-AUC) evaluated the classifier's ability to rank positive instances higher than negative ones across varying classification thresholds.

The models were evaluated against an unseen testing dataset consisting of 56,962 transactions, which included 98 actual frauds. The baseline Logistic Regression model utilizing class weights but no synthetic data successfully captured a high percentage of fraud with a recall of 0.92, but exhibited a severely degraded precision score of 0.06. Applying SMOTE to the Logistic Regression model improved the decision boundary slightly, yielding a precision of 0.13 and a recall of 0.90, alongside a theoretical ROC-AUC of 0.976. However, the linear algorithm still struggled to filter out false positives effectively. The Random Forest ensemble definitively emerged as the optimal architecture, achieving a precision of 0.84, a recall of 0.83, an F1-score of 0.83, and an ROC-AUC of 0.964.

The operational efficacy of the proposed ensemble system is best observed through its absolute prediction counts. As illustrated in Figure 8: Confusion Matrix, the Random Forest model achieved 56,848 true negatives, 81 true positives, 17 false negatives, and a mere 16 false positives. In a production banking environment, the cost matrix is dual-faceted. The baseline linear model generated 1,386 false positives on the exact same test set, representing an operationally unviable volume that would result in severe consumer friction and massive customer service overhead. In stark contrast, the Random Forest model's minimal false positive rate guarantees virtually zero disruption to legitimate consumers while simultaneously capturing the vast majority of malicious activity.

A primary advantage of the Random Forest architecture is its inherent interpretability. By calculating the mean decrease in Gini impurity across the ensemble of trees, the model quantifiably isolates the most critical predictive features. The analysis identified the top principal components contributing to the fraud classification vectors, specifically isolating V14, V10, V4, V12, and V17 as the most influential variables. These components represent the most distinct statistical deviations between legitimate user behavior and compromised credit card activity, serving as critical focal points for future heuristic rule synthesis, as visualized in Figure 7: Feature Importance.

Table 1: Model Performance Comparison

Model Architecture	SMOTE Applied	Precision	Recall	F1-Score	ROC-AUC
Logistic Regression	No	0.06	0.92	0.11	0.972
Logistic Regression	Yes	0.13	0.90	0.23	0.976
Random Forest	**Yes**	**0.84**	**0.83**	**0.83**	**0.964**

IV. CONCLUSION

This research successfully engineered and evaluated a highly resilient machine learning pipeline for credit card fraud detection, directly addressing the algorithmic vulnerabilities associated with extreme class imbalance. The empirical exploration confirms that training predictive algorithms on raw, imbalanced financial data yields heavily biased systems that fall victim to the accuracy paradox. However, through the strategic, localized application of the Synthetic Minority Over-

ISSN (online): 2581-3048
Volume 10, Issue 4, pp 61-65, April-2026
https://doi.org/10.47001/IRJIET/2026.104007

sampling Technique (SMOTE) strictly within the training environment, the feature space was successfully calibrated. This methodology definitively prevented data leakage, ensuring that the resulting evaluation metrics reflect true, real-world predictive capability rather than synthetically inflated performance.

The comparative analysis yielded a clear, consequential outcome: while baseline linear models like Logistic Regression demonstrated high theoretical class separation, their severely degraded precision rendered them highly susceptible to false alarms. The Random Forest Classifier definitively emerged as the superior architecture. By achieving an optimal balance between precision (0.84) and recall (0.83), the ensemble model successfully minimized false negative rates without inflating false positive rates—generating a negligible 16 false positives out of 56,864 transactions. The significance of this outcome is profound for institutional deployment; it provides payment gateways with a robust, automated defense mechanism that protects institutional capital by intercepting malicious activity, while simultaneously preserving a frictionless, undisrupted experience for legitimate consumers.

Looking forward, the consequences of this research provide a strong foundation for next-generation financial security systems. While the current architecture demonstrates exceptional performance on batch-processed data, future iterations will focus on deploying the model within a real-time streaming environment utilizing distributed event streaming platforms. Additionally, integrating cost-sensitive learning algorithms—which assign heavier mathematical penalties to misclassified fraudulent transactions—and exploring deep learning paradigms such as Artificial Neural Networks will further enhance the system's ability to autonomously adapt to the evolving, complex concept drift of modern cyber-financial threats.

ACKNOWLEDGEMENT

The authors express their profound gratitude to the Department of Artificial Intelligence and Machine Learning (AIML) for providing the computational resources and administrative support necessary to conduct this algorithmic research. We also extend our sincere appreciation to our project guide and faculty members for their invaluable academic expertise, continuous mentorship, and constructive feedback throughout the development and evaluation of this comparative analysis framework.

REFERENCES

[1] R. Bin Sulaiman, V. Schetinin & P. Sant, Review of Machine Learning Approach on Credit Card Fraud Detection, *Human-Centric Intelligent Systems*, 2022. — comprehensive ML methods & challenges.

[2] K. Ghosh Dastidar, O. Caelen & M. Granitzer, Machine Learning Methods for Credit Card Fraud Detection: A Survey, *IEEE Access*, 2024.

[3] Y. A. Hassan & O. S. Kareem, Credit Card Fraud Detection: Comparative Study of ML & DL Methods, *Engineering and Technology Journal*, 2025. — recent ML vs deep learning comparison.

[4] E. Btoush et al., Achieving Excellence in Cyber Fraud Detection: A Hybrid ML+DL Ensemble Approach for Credit Cards, *Applied Sciences*, 2025. — ensemble and hybrid ML+DL strategies.

[5] E. Ileberi, Y. Sun & Z. Wang, ML-based Credit Card Fraud Detection with GA Feature Selection, *Journal of Big Data*, 2022 — uses feature selection + multiple classifiers.

[6] E. Btoush et al., Resampling Methods for Imbalanced Credit Card Fraud Data, *Applied Sciences*, 2026 — analysis of sampling + ML classifiers (XGBoost, RF, etc.).

[7] Autonomous credit card fraud detection using LSTM-RNN, *Computers & Electrical Engineering*, 2022 — RNN + deep learning methods.

[8] Frontiers in AI, Enhancing Credit Card Fraud Detection using Traditional and Deep Learning Models with Imbalance Mitigation, 2025 — modern deep learning evaluation.

[9] R. Jayalakshmi, R. G. S. Kumar & T. Thanushree, A Survey on Credit Card Fraud Detection Using Deep Learning Models, *IRJAEM*, 2025 — DL-focused survey.

[10] R. Ali, Explainable AI Framework for Credit Card Fraud Detection using XGBoost & Deep Neural Networks, *SSRN*, 2026 — blends ML, DL & SHAP for interpretability.

AUTHORS BIOGRAPHY

Prathmesh is an AI & ML student at Ajeenkya D.Y. Patil School of Engineering, researching predictive analytics and machine learning for financial anomaly detection.

Anjaneya is an AI & ML student at Ajeenkya D.Y. Patil School of Engineering, specializing in data preprocessing and handling imbalanced datasets with SMOTE.

International Research Journal of Innovations in Engineering and Technology (IRJIET)

ISSN (online): 2581-3048

Volume 10, Issue 4, pp 61-65, April-2026

https://doi.org/10.47001/IRJIET/2026.104007

Suyash is an AI & ML student at Ajeenkya D.Y. Patil School of Engineering, focusing on ensemble machine learning architectures and cybersecurity applications.

Citation of this Article:

Prathmesh Sunil Dhobe, Anjaneya Kokre, Suyash Kale, Kartik Sabe, & Shahrukh Shaikh. (2026). Credit Card Fraud Detection: Mitigating Extreme Class Imbalance Using Synthetic Oversampling and Ensemble Machine Learning. *International Research Journal of Innovations in Engineering and Technology - IRJIET*, 10(4), 61-65. Article DOI https://doi.org/10.47001/IRJIET/2026.104007

****

