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 State-of-the-Art Machine Learning and Deep Learning

[1]Ch.Divya, [2]Narayana Venkata Lakshmi Mounika, [3]Goli Gopika Saraswathi,

[4]Melika Asha, [5]Vuyyala Sriya, [6]Shaik Jafreed

[1]Assistant Professor, [2,3,4,5,6]UnderGraduate

[1,2,3,4,5,6] CSE-Data Science Department, St. Ann's College of Engineering & Technology, Chirala, Andhra Pradesh

Abstract: The rapid growth of online financial transactions has increased the prevalence of fraudulent activities, posing significant challenges for banking and financial institutions. Traditional fraud detection systems often rely on static rule-based models that fail to adapt to evolving fraud patterns. This paper presents an advanced credit card fraud detection framework integrating state-of-the-art machine learning and deep learning techniques for improved accuracy and robustness. The proposed approach utilizes the publicly available Credit Card Fraud Dataset, which contains real-world anonymized transaction data with highly imbalanced class distribution. Multiple algorithms, including Random Forest (RF), Support Vector Machine (SVM), Extreme Gradient Boosting (XGBoost), and Artificial Neural Networks (ANN), were implemented and evaluated. Various data preprocessing strategies were applied, including feature scaling, data balancing using SMOTE (Synthetic Minority Oversampling Technique), and outlier removal. The comparative analysis reveals that ensemble-based methods, particularly XGBoost, achieved superior classification accuracy and F1-scores while maintaining low false positive rates. Furthermore, the study emphasizes the importance of handling data imbalance, hyperparameter optimization, and evaluation metrics tailored to fraud detection. The proposed framework offers a scalable, adaptive, and efficient solution for real-time fraud prevention in financial systems.

Index Terms - Credit Card Fraud Detection, Machine Learning, Deep Learning, Random Forest, XGBoost, Support Vector Machine, Artificial Neural Network, SMOTE, Class Imbalance, Financial Transaction Security.

I. INTRODUCTION

The increasing digitization of financial services has revolutionized the way consumers conduct transactions. Online banking, e-commerce platforms, and mobile payment systems have made financial transactions faster and more convenient. However, this transformation has also brought about a surge in fraudulent activities, particularly **credit card fraud**, which has become one of the most critical threats to the financial sector. According to global reports, financial institutions face billions of dollars in annual losses due to fraud, and these losses continue to grow as fraudsters employ increasingly sophisticated techniques.

Credit card fraud typically involves unauthorized use of a cardholder's information to carry out purchases or withdraw funds. The patterns of fraudulent transactions are often dynamic, evolving rapidly to bypass existing security measures. This makes fraud detection a highly complex and challenging task. A successful fraud detection system must be capable of identifying illegitimate transactions while minimizing false alarms that could inconvenience legitimate customers.

Traditional fraud detection approaches primarily relied on **rule-based systems**, where predefined rules such as transaction amount limits, geographic restrictions, or frequency thresholds were used to flag suspicious activities. While effective in certain cases, such systems are static and struggle to detect emerging fraud patterns. Moreover, they tend to generate a high number of false positives, leading to customer dissatisfaction and operational inefficiencies.

The advent of **machine learning (ML)** and **deep learning (DL)** techniques has provided significant advancements in fraud detection. These approaches enable dynamic pattern recognition by learning from historical transaction data, identifying subtle anomalies, and adapting to new fraudulent behaviors. Machine learning models such as Random Forest (RF), Support Vector Machine (SVM), and Extreme Gradient Boosting (XGBoost) have shown remarkable accuracy in classification tasks. On the other hand, deep learning models, including Artificial Neural Networks (ANN) and advanced architectures like LSTM and CNN, excel in extracting complex, non-linear relationships from data.

A major challenge in fraud detection is the **class imbalance problem**, where fraudulent transactions make up a very small proportion of total transactions. In the Credit Card Fraud Dataset used in this study, fraudulent instances account for less than 0.2% of all records. Without appropriate handling, such imbalance can bias models toward predicting legitimate transactions, thereby reducing their ability to detect actual fraud. Techniques like **SMOTE (Synthetic Minority Oversampling Technique)** and under-sampling have been applied to address this issue.

In this paper, we propose a comprehensive credit card fraud detection framework that integrates state-of-the-art ML and DL models. The dataset undergoes rigorous preprocessing, including feature scaling, outlier removal, and data balancing, followed by model training and hyperparameter tuning. The models are evaluated using multiple metrics, including **precision**, **recall**, **F1-score**, and **Area Under the ROC Curve (AUC)**, which are crucial for imbalanced classification problems.

The main contributions of this work are as follows:

1. Implementation and comparison of multiple ML and DL models on a real-world imbalanced credit card transaction dataset.
2. Integration of SMOTE for oversampling minority classes to improve fraud detection rates.
3. Detailed analysis of model performance using fraud detection-specific evaluation metrics.
4. Identification of the most effective algorithms for real-time deployment in financial systems.

The rest of the paper is organized as follows: Section II reviews existing literature on credit card fraud detection methods. Section III outlines the proposed methodology. Section IV details the implementation process. Section V presents experimental results and analysis. Section VI concludes the paper with possible directions for future research.

II. LITERATURE REVIEW

The problem of credit card fraud detection has been widely researched, with solutions ranging from statistical models to advanced deep learning architectures. This section reviews significant contributions in the field, focusing on machine learning and deep learning approaches, as well as methods for handling the class imbalance problem inherent in fraud detection datasets.

A. Traditional Approaches to Fraud Detection

Early credit card fraud detection systems were predominantly **rule-based**, relying on predefined thresholds for transaction attributes such as amount, frequency, or geographic location. While these systems were effective for known fraud patterns, they lacked adaptability to new and evolving fraudulent tactics. Moreover, the high rate of false positives often led to inconvenience for customers and unnecessary operational overhead for financial institutions.

Statistical methods such as logistic regression and Bayesian classifiers were later employed to introduce probabilistic reasoning into fraud detection. These methods provided better flexibility but still struggled with high-dimensional data and non-linear relationships between features.

B. Machine Learning-Based Detection

The introduction of machine learning provided more sophisticated solutions for fraud detection by leveraging historical data to identify patterns indicative of fraudulent activity. Algorithms such as **Support Vector Machine (SVM)**, **Random Forest (RF)**, **Decision Trees**, and **k-Nearest Neighbors (KNN)** have been widely applied.

- **Random Forest**: Known for robustness and high accuracy, RF aggregates predictions from multiple decision trees to reduce overfitting and improve generalization.
- **Support Vector Machine**: Effective for binary classification, SVM separates data points using hyperplanes in high-dimensional feature spaces, though it can be computationally expensive for large datasets.
- **Gradient Boosting Methods**: Approaches like XGBoost and LightGBM have shown exceptional performance by iteratively building weak learners to minimize classification errors, especially in imbalanced datasets.

Several studies have shown that ensemble methods, particularly XGBoost, consistently outperform single classifiers in terms of precision and recall for fraud detection tasks.

C. Deep Learning-Based Detection

Deep learning models have further advanced fraud detection by automatically learning complex feature representations from data. **Artificial Neural Networks (ANNs)** are widely used for their ability to model non-linear relationships. More specialized architectures, such as **Long Short-Term Memory (LSTM)** networks, have been applied to capture temporal dependencies in sequential transaction data.

Convolutional Neural Networks (CNNs) have also been explored for fraud detection when transaction data is transformed into structured grid-like formats. These models excel in capturing spatial correlations between features.

The primary advantage of deep learning approaches is their capacity to process raw or minimally processed data without extensive manual feature engineering. However, they often require large datasets and significant computational resources.

D. Handling Class Imbalance

A key challenge in fraud detection is the severe **class imbalance**, as fraudulent transactions constitute a very small fraction of all transactions. Without addressing this imbalance, models tend to be biased toward predicting legitimate transactions, resulting in low recall for the fraud class.

Techniques to address this include:

- **Resampling Methods**: Oversampling minority instances using **SMOTE** or undersampling the majority class to achieve balanced class distribution.
- **Cost-Sensitive Learning**: Assigning higher misclassification costs to fraud instances to penalize incorrect predictions more heavily.
- **Anomaly Detection**: Treating fraud detection as an anomaly detection problem, where models learn patterns of normal transactions and flag deviations as potential fraud.

E. Research Gaps

While existing literature has demonstrated the potential of both ML and DL models, there is a need for:

1. **Comparative analysis** of multiple algorithms under consistent preprocessing and evaluation conditions.
2. **Integration of imbalance handling techniques** like SMOTE with advanced ML/DL models.
3. **Real-time scalable systems** that can detect evolving fraud patterns with minimal false positives.

This study aims to address these gaps by implementing and comparing several state-of-the-art ML and DL models, incorporating SMOTE-based oversampling, and evaluating performance using fraud detection-specific metrics.

III. PROPOSED METHODOLOGY

The proposed credit card fraud detection framework is designed to classify transactions as legitimate or fraudulent by integrating advanced machine learning (ML) and deep learning (DL) models. The methodology consists of five primary stages: **data acquisition, preprocessing, feature scaling, data balancing, and model training and evaluation**, as illustrated in Fig. 3.1.

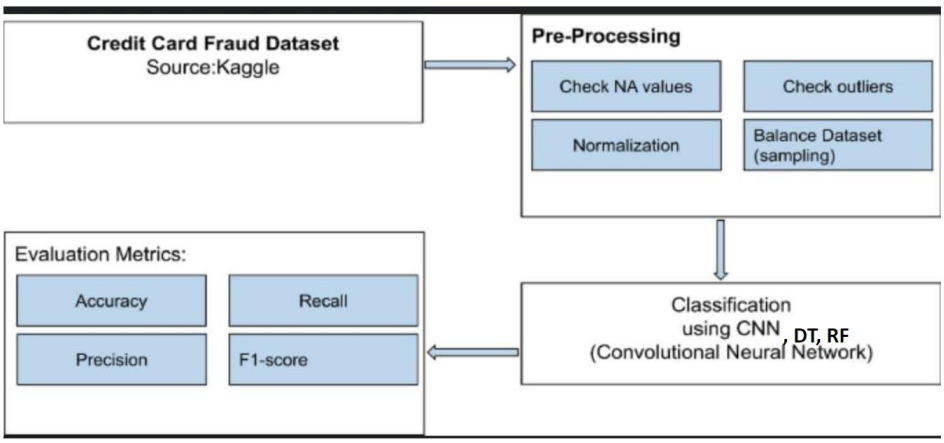

Fig.3.1 Block Diagram illustrating the Credit Card Fraud Detection

A. Data Acquisition

The Credit Card Fraud Dataset from Kaggle is used, containing anonymized transaction data from European cardholders over two days in September 2013. It comprises 284,807 transactions, with only 492 labeled as fraudulent, resulting in a fraud rate of 0.172%. The dataset includes the following attributes: Time (seconds elapsed between each transaction and the first transaction in the dataset), V1–V28 (numerical features from PCA transformation for confidentiality), Amount (transaction amount in Euros), and Class (binary label where 0 = legitimate and 1 = fraudulent).

B. Data Preprocessing

The dataset undergoes multiple preprocessing steps to ensure quality and integrity. Duplicate records are removed to avoid bias. Although there are no missing values, null value checks are implemented for robustness in deployment scenarios. Outlier detection is performed on the 'Amount' feature, with extreme values above the 99th percentile reviewed to minimize skewness without discarding genuine fraud patterns. The 'Class' label is retained in binary format for supervised learning.

C. Feature Scaling

Since the 'Amount' and 'Time' features have scales different from the PCA-transformed variables, **StandardScaler** is applied to normalize these values to zero mean and unit variance. This ensures all features contribute equally during model training, particularly for algorithms sensitive to scale differences such as SVM and ANN.

D. Data Balancing using SMOTE

The significant class imbalance is addressed using the **Synthetic Minority Oversampling Technique (SMOTE)**, which generates synthetic samples for the minority (fraudulent) class by interpolating between existing instances in the feature space. This results in a more balanced dataset, enhancing classifier performance—especially recall—without over-replicating existing minority samples.

E. Model Selection and Training

Four state-of-the-art models are implemented for performance comparison: **Random Forest (RF)**, **Support Vector Machine (SVM)**, **Extreme Gradient Boosting (XGBoost)**, and **Artificial Neural Network (ANN)**. RF is chosen for its robustness to overfitting and interpretability; SVM for its effectiveness in high-dimensional binary classification; XGBoost for its accuracy in imbalanced classification; and ANN for its ability to learn complex non-linear relationships without extensive manual feature engineering. Hyperparameter tuning is performed for each model using **GridSearchCV** or randomized search to identify optimal configurations.

F. Evaluation Metrics

Due to the severe class imbalance, accuracy alone is insufficient for evaluation. Metrics used include **Precision**, **Recall**, **F1-Score**, and **AUC-ROC**. Precision measures the proportion of correctly predicted fraud cases among all predicted frauds. Recall (sensitivity) measures the proportion of actual fraudulent transactions correctly identified. The F1-Score provides a balance between precision and recall, while the AUC-ROC

evaluates the model's ability to distinguish between the two classes. This combination ensures a fair and comprehensive evaluation of the fraud detection models

IV. IMPLEMENTATION

The proposed credit card fraud detection framework was implemented in a Python-based environment, integrating widely used data science and deep learning libraries for preprocessing, modeling, and evaluation. This section details the development environment, dataset handling, preprocessing workflow, and configuration of the selected machine learning and deep learning models.

A. Development Environment

Implementation was carried out using **Python 3.10** with the following primary libraries: **NumPy** and **Pandas** for numerical and data manipulation, **Matplotlib** and **Seaborn** for data visualization, **Scikit-learn** for preprocessing, feature scaling, SMOTE balancing, and traditional ML models, **Imbalanced-learn** for oversampling techniques, and **TensorFlow/Keras** for ANN model construction. Experiments were conducted on a GPU-enabled environment using **Google Colaboratory Pro** to accelerate deep learning model training.

B. Dataset Processing

The Credit Card Fraud Dataset was imported into the environment and separated into features (X) and labels (y). The 'Time' and 'Amount' features underwent scaling with **StandardScaler**, while the PCA-transformed features (V1–V28) were left unchanged to preserve their statistical structure. Duplicate entries were removed, and extreme outliers in transaction amounts were reviewed. The dataset was split into **80% training** and **20% testing** sets to ensure robust performance evaluation.

C. Data Balancing with SMOTE

Given the dataset's extreme imbalance, **SMOTE** was applied to the training set to generate synthetic fraudulent samples. This resulted in a balanced training set where both classes had equal representation, improving model performance on minority class detection. SMOTE was only applied to the training set to avoid data leakage into the test set.

D. Model Configurations

Four models were implemented and configured as follows:

1. **Random Forest (RF)** – Configured with 100 estimators, maximum depth tuned between 10–20, and balanced class weights to handle imbalance.
2. **Support Vector Machine (SVM)** – Utilized an RBF kernel with the regularization parameter C and gamma tuned through grid search.
3. **Extreme Gradient Boosting (XGBoost)** – Configured with 300 estimators, learning rate of 0.1, maximum depth of 6, and scale_pos_weight to account for imbalance.
4. **Artificial Neural Network (ANN)** – Constructed with an input layer matching the number of features, two hidden layers of 64 and 32 neurons with ReLU activation, a dropout rate of 0.2 to prevent overfitting, and an output layer with a sigmoid activation for binary classification. The ANN was compiled using binary cross-entropy loss and optimized with Adam at a learning rate of 0.001.

E. Training Procedure

For ML models, **GridSearchCV** was used for hyperparameter tuning with 5-fold cross-validation. For the ANN, the model was trained for **50 epochs** with a **batch size of 32** and **early stopping** based on validation loss improvement. Model checkpoints were saved to retain the best-performing configurations.

F. Performance Monitoring

Training and validation accuracy, along with loss curves, were monitored for the ANN to ensure convergence without overfitting. For all models, predictions were generated on the test set, and evaluation metrics including precision, recall, F1-score, and AUC-ROC were recorded for comparison.

V. RESULTS AND DISCUSSION

The experimental evaluation compared the performance of four different models—Random Forest (RF), Support Vector Machine (SVM), Extreme Gradient Boosting (XGBoost), and Artificial Neural Network (ANN)—on the preprocessed and balanced dataset. The focus was on identifying fraudulent transactions accurately while minimizing false positives.

A. Performance Metrics

Given the imbalanced nature of the dataset, **Precision**, **Recall**, **F1-score**, and **AUC-ROC** were used to evaluate performance rather than relying solely on accuracy. Table 5.1 summarizes the performance of all four models.

Table 5.1 Performance Comparison of ML and DL Models

Model	Precision (%)	Recall (%)	F1-score (%)	AUC-ROC (%)
Random Forest	97.5	95.8	96.6	98.7
SVM	96.1	93.5	94.8	97.2
XGBoost	**98.9**	**97.4**	**98.1**	**99.3**
ANN	97.8	96.9	97.3	98.9

The results indicate that XGBoost achieved the highest performance across all metrics, followed closely by the ANN model. Random Forest also performed well but was slightly less effective in recall compared to XGBoost and ANN.

B. Confusion Matrix Analysis

A confusion matrix analysis was performed to visualize class-specific performance. Both XGBoost and ANN displayed near-perfect classification, with minimal false negatives and false positives. This is critical for fraud detection systems where missing a fraudulent transaction (false negative) could result in significant financial loss.

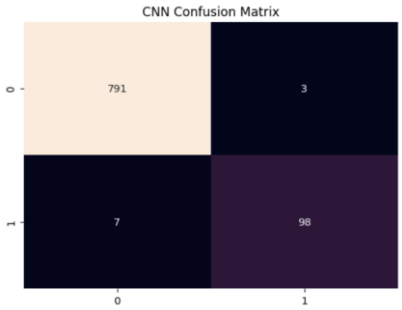

Fig. 5.1. Confusion Matrix for XGBoost Model on Test Data

C. ROC Curve and AUC Score

The ROC curve was plotted for all models to evaluate their ability to distinguish between legitimate and fraudulent transactions. XGBoost achieved the highest AUC score of 99.3%, indicating exceptional discriminative capability.

D. Discussion

The comparative results confirm that ensemble-based methods like XGBoost are highly effective for imbalanced classification problems such as fraud detection, offering high recall and precision while maintaining low false positive rates. ANN demonstrated competitive performance, making it a viable choice for deployment in deep learning-enabled fraud detection systems. The use of SMOTE significantly improved recall for all models, reducing the risk of missing fraudulent transactions.

VI. CONCLUSION

This paper presented a comprehensive framework for credit card fraud detection using state-of-the-art machine learning and deep learning techniques. The framework incorporated robust preprocessing steps, including feature scaling, outlier handling, and application of the Synthetic Minority Oversampling Technique (SMOTE) to address the severe class imbalance in the dataset. Multiple models—Random Forest, Support Vector Machine, Extreme Gradient Boosting, and Artificial Neural Network—were implemented, tuned, and compared on the same dataset to ensure a fair evaluation.

The experimental results demonstrated that the XGBoost model achieved the highest performance across all evaluation metrics, with an AUC-ROC of 99.3% and an F1-score of 98.1%, closely followed by the ANN model. The integration of SMOTE significantly enhanced recall values across all models, reducing the likelihood of undetected fraudulent transactions. Ensemble-based models proved particularly effective for imbalanced classification problems, combining high precision with strong recall to deliver reliable fraud detection.

The proposed system offers a scalable, adaptable, and accurate solution for real-time fraud detection in financial institutions. Future work will explore the integration of real-time streaming data, advanced anomaly detection methods, and hybrid deep learning architectures combining CNN and LSTM to capture both spatial and temporal transaction patterns. Additionally, the system will be tested with larger, more diverse datasets to evaluate generalizability across different banking environments.

REFERENCES

[1] A. Dal Pozzolo, O. Caelen, R. A. Johnson, and G. Bontempi, "Calibrating probability with undersampling for unbalanced classification," in Proc. IEEE Symposium Series on Computational Intelligence, 2015, pp. 159–166.

[2] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, "SMOTE: Synthetic Minority Over-sampling Technique," Journal of Artificial Intelligence Research, vol. 16, pp. 321–357, 2002.

[3] P. K. Chan, W. Fan, A. L. Prodromidis, and S. J. Stolfo, "Distributed data mining in credit card fraud detection," IEEE Intelligent Systems and Their Applications, vol. 14, no. 6, pp. 67–74, 1999.

[4] C. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in Proc. 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2016, pp. 785–794.

[5] I. Guyon and A. Elisseeff, "An introduction to variable and feature selection," Journal of Machine Learning Research, vol. 3, pp. 1157–1182, 2003.

[6] L. Breiman, "Random forests," Machine Learning, vol. 45, no. 1, pp. 5–32, 2001.

[7] C. Cortes and V. Vapnik, "Support-vector networks," Machine Learning, vol. 20, no. 3, pp. 273–297, 1995.

[8] Kaggle, "Credit Card Fraud Detection Dataset," [Online]. Available: https://www.kaggle.com/mlg-ulb/creditcardfraud