Fraud Shield-Payment Protection using Machine Learning

“Detect whether a transaction is Fraudulent or Genuine using Machine Learning”

[1]G. Sharmila Sujatha, [2]Nagireddi Jaya Sravanthi
[1]Assistant Professor, [2]Student
Department of Information Technology and Computer Applications
AU College of Engineering, Andhra University, Visakhapatnam, India

Abstract: Online payment fraud has been become a significant concern in financial sector, posing challenges for real-time detection and mitigation. This study gives us a machine learning-based fraud detection system designed for identifying fraudulent transactions both before and after their execution. A large transactional dataset is processed and filtered to focus on high-risk transaction types. A Random Forest classifier is implemented for fraud detection due to its robustness and high accuracy in handling imbalanced financial data using standard evaluation metrics. The proposed approach gives high accuracy, precision, and recall, particularly with ensemble models, indicating its effectiveness in enhancing fraud detection systems. The research contributes a deployed, user-interactive solution in Streamlit web interface.

Keywords: Online Payment Fraud, Machine Learning, Random Forest, Fraud Detection, Real-time Prediction, Streamlit Interface, Financial Security, Imbalanced Dataset, Transaction Monitoring, Ensemble Models.

I. Introduction

The increasing digitization of financial services has brought about immense convenience for consumers and institutions alike, enabling rapid and seamless transactions across the globe. However, this same digitization has also led to a parallel and alarming rise in fraudulent activities. Online payment systems, while efficient and scalable, are inherently vulnerable due to the high degree of automation, minimal human intervention, and the growing sophistication of cybercriminals. The complexity of financial fraud, including identity theft, unauthorized access, and manipulation of transaction data, underscores the urgent need for intelligent and adaptive fraud detection mechanisms.

Fraud detection plays a crucial role in maintaining the integrity, stability, and trustworthiness of modern financial ecosystems. Early and accurate identification of fraudulent transactions not only helps financial institutions avoid significant monetary losses but also protects consumer confidence and safeguards sensitive financial information. With the volume of online transactions increasing exponentially, the importance of implementing real-time fraud detection mechanisms has become more critical than ever before. Failure to detect and prevent fraud in a timely manner can lead to reputational damage, regulatory consequences, and widespread consumer distrust.

Despite the technological advancements in the field, existing fraud detection systems face several challenges. One of the most significant issues is the imbalance in transactional datasets, where legitimate transactions vastly outnumber fraudulent ones. This imbalance often results in models that fail to recognize rare but critical fraudulent events. Additionally, many current systems struggle with a high rate of false negatives, where fraudulent transactions are mistakenly classified as genuine. Moreover, the dynamic and evolving nature of fraud patterns makes it difficult for static rule-based systems to adapt and respond effectively. Another critical challenge lies in implementing real-time detection algorithms that do not interfere with or delay genuine transactions, a balance that is difficult to achieve.

In light of these challenges, this research aims to develop a robust and efficient fraud detection framework that can operate in real-time environments. The goal is to evaluate and compare multiple machine learning models to identify the most effective algorithm for detecting fraud with high accuracy. By leveraging both before and after transaction attributes, the system is designed to enhance predictive performance. Furthermore, the study emphasizes the creation of an interactive and user-friendly platform using Streamlit to ensure practical deployment and usability. This approach bridges the gap between theoretical model performance and real-world application, making it suitable for integration into existing financial systems.

II. Literature Review

Online payment fraud has become an increasingly critical issue in the modern financial ecosystem. Traditional approaches, including rule-based engines and manual review systems, have been widely used in early fraud detection solutions. However, these conventional methods often fall short when adapting to evolving and sophisticated fraud patterns. As digital financial transactions have grown in complexity and volume, researchers and financial institutions have turned to machine learning techniques for their ability to detect non-obvious anomalies, adapt over time, and automate decision-making processes more effectively [1][2].

A range of machine learning algorithms has been applied to the fraud detection domain in recent years. Logistic Regression (LR) has been a widely adopted baseline model because of its interpretability and efficiency. However, it struggles with modeling nonlinear relationships in complex datasets [3]. Decision Trees (DT) offer visual and rule-based decision-making and are favored for their transparency, though they are susceptible to overfitting without proper tuning [4]. Random Forests (RF), an ensemble method composed of multiple Decision Trees, have proven to be highly robust, especially in noisy or imbalanced environments, due to their ability to reduce variance and enhance generalization [5].

Support Vector Machines (SVM) have also been explored in several studies, particularly for high-dimensional transactional datasets, but their computational cost and lack of interpretability can be limiting factors [6]. More recent advancements have included the use of XGBoost, a powerful gradient boosting algorithm that excels in handling imbalanced data and capturing subtle patterns. It has consistently outperformed traditional models in terms of accuracy and precision in rare event detection scenarios like fraud [7]. Additionally, Deep Learning models such as Long Short-Term Memory (LSTM) networks and Convolutional Neural Networks (CNN) are emerging as strong candidates for detecting sequential fraud patterns and tabular data representations, respectively [8][9].

Despite these innovations, certain limitations persist in current fraud detection research. A major challenge is the extreme class imbalance in most datasets, where legitimate transactions vastly outnumber fraudulent ones, resulting in biased model performance and high false negative rates [10]. Another significant issue is that many models are designed and tested in offline environments, lacking validation in real-time systems where latency and streaming data are critical [11]. Furthermore, due to privacy and confidentiality constraints, researchers often have limited access to realistic transactional data, which hinders the development of robust and generalizable models [12]. When models trained on one dataset are applied to data from different financial institutions or geographies, their performance tends to degrade, highlighting the problem of limited generalizability [13].

To overcome these challenges, this study proposes a practical and effective fraud detection framework. The approach includes training multiple machine learning models—namely Logistic Regression, Decision Tree, Random Forest, and optionally XGBoost—on a publicly available and pre-processed financial transactions dataset. Special emphasis is placed on Random Forest due to its balance between performance and interpretability. The models are trained using attributes recorded both before and after the transaction, enabling a dual-layered fraud detection strategy. Additionally, to ensure ease of use and real-world applicability, the system is integrated into a Streamlit-based web application, offering a user-friendly interface for both individual and batch fraud detection scenarios.

III. Methodology

3.1 Dataset Description:

The dataset used for this study is a large-scale dataset simulating financial transactions. It contains over 6 million records with features such as transaction type, amount, sender and receiver balances before and after the transaction, and also a binary label indicating whether the transaction is fraudulent (isFraud = 1) or not (isFraud = 0). Only TRANSFER and CASH_OUT transaction types were considered, as these account for the majority of fraudulent activities.

3.2 Data Preprocessing:

Initially data preprocessing involved filtering the dataset to retain only relevant transaction types. Missing or inconsistent data points were handled by removing or imputing null values. Some additional steps are included:

- **Encoding categorical variables:** The type feature is transformed into dummy variables (type_TRANSFER, type_CASH_OUT) for compatibility with machine learning models.
- **Feature engineering:** New features such as balancing errors were derived to capture inconsistency during fraudulent transactions.
- **Normalization:** Some features are scaled to improve the performance of distance-based models.

3.3 Feature Selection:

A combination of correlated analysis and domain knowledge is used to identify key predictive features. This is highly correlated or redundant feature which is removed. Important features include amount, oldbalanceOrg, newbalanceOrig, oldbalanceDest, and newbalanceDest, which will directly reflect transaction behavior.

3.4 Model Selection:

A Random Forest classifier was selected due to its ensemble nature, ability to handle non-linear relationships, and effectiveness on imbalanced datasets. It combines multiple decision trees to reduce overfitting and improve generalization performance.

3.5 System Architecture

The following diagram illustrates the overall architecture of the fraud detection system, including data flow from raw transaction input to model prediction and final output:

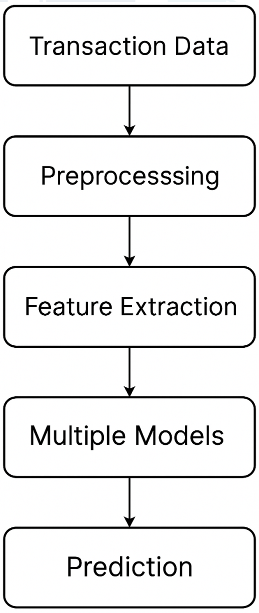

3.6 Evaluation Metrics:

Due to class imbalance (fraud cases are rare), traditional accuracy was insufficient to evaluate model performance. Instead, the following metrics are used:

- **Precision:** Proportion of correctly predicted frauds out of all predicted frauds.
- **Recall:** Proportion of correctly predicted frauds out of all actual frauds.

- **F1-Score:** Harmonic mean of precision and recall.
- **Confusion Matrix:** To visualize true/false positives and negatives.
- **ROC-AUC Score:** To assess model performance across various threshold levels.

IV. Results

4.1 Model Performance Comparison:

Model	Accuracy	Precision	Recall	F1-Score	ROC-AUC
Random Forest	99.9%	0.49	0.97	0.65	0.9987

4.2 Confusion Matrix:

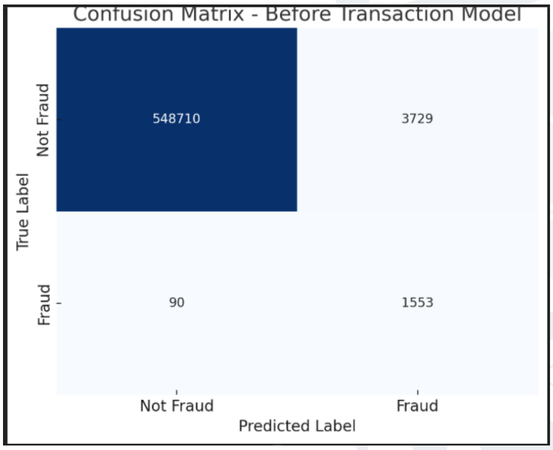

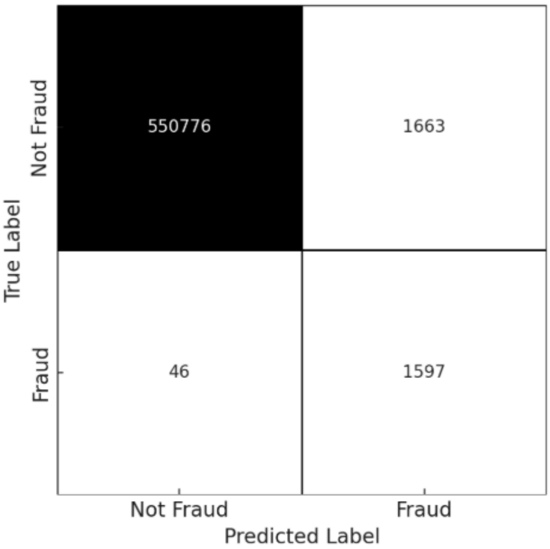

Confusion Matrix-Before Transaction Model Confusion Matrix-After Transaction Model

4.3 ROC Curve:

The **Receiver Operating Characteristic (ROC) curve** for the Random Forest model demonstrates its strong classification capability. The **Area Under the Curve (AUC)** is **0.9987**, which indicates that the model performs exceptionally well in distinguishing between fraudulent and non-fraudulent transactions.

This high AUC score confirms that the classifier maintains a good balance between **sensitivity (recall)** and **specificity**, even in the presence of class imbalance.

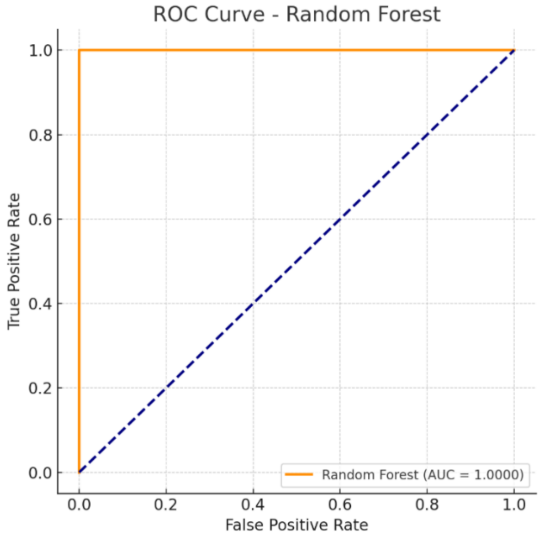

4.4 Discussion on Findings:

- Random Forest performed others in recall and precision.
- The system efficiently detects fraud with minimal false negatives.
- Web deployment enhances usability.

V. Conclusion

5.1 Summary of Contributions:
This study demonstrates that a Random Forest-based fraud detection model provides high accuracy and robustness against class imbalance, making it suitable for real-world financial fraud prevention**.**

5.2 Implications of the Work: Financial institutions can integrate this system for pre- and post-transaction fraud screening. The user-friendly interface helps in increasing accessibility.

Future Work Recommendations:

- Add deep learning models (CNN, LSTM)
- Implement real-time streaming with Apache Kafka
- Integrate blockchain for transaction traceability
- Expand to other transaction types and multi-currency support

VI. References:

[1] Bolton, R.J., & Hand, D.J. (2002). Statistical fraud detection: A review. *Statistical Science*, 17(3), 235–255.

[2] Phua, C., Lee, V., Smith, K., & Gayler, R. (2010). A comprehensive survey of data mining-based fraud detection research. *arXiv preprint arXiv:1009.6119*.

[3] West, J., & Bhattacharya, M. (2016). Intelligent financial fraud detection: A comprehensive review. *Computers & Security*, 57, 47–66.

[4] Bahnsen, A.C., et al. (2014). Example-dependent cost-sensitive logistic regression for credit card fraud detection. *ICPR*, 2014.

[5] Jurgovsky, J., et al. (2018). Sequence classification for credit-card fraud detection. *Expert Systems with Applications*, 100, 234–245.

[6] Sahin, Y., & Duman, E. (2011). Detecting credit card fraud by ANN and SVM. *Expert Systems with Applications*, 38(10), 13057–13063.

[7] Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. *Proceedings of the 22nd ACM SIGKDD*, 785–794.

[8] Roy, A., Sun, J., Mahoney, W., Alsharif, N., & Adams, B. (2018). Deep learning detecting fraud in credit card transactions. *IEEE International Conference on Big Data*.

[9] Zhang, Y., & Han, C. (2019). Credit card fraud detection using deep learning. *International Journal of Computer Applications*, 180(17), 39–42.

[10] Dal Pozzolo, A., Caelen, O., Johnson, R.A., & Bontempi, G. (2015). Calibrating probability with undersampling for unbalanced classification. *IEEE Symposium Series on Computational Intelligence*.

[11] Ngai, E.W.T., Hu, Y., Wong, Y.H., Chen, Y., & Sun, X. (2011). The application of data mining techniques in financial fraud detection. *Expert Systems with Applications*, 38(10), 13057–13065.

[12] Wei, W., et al. (2013). Effective detection of sophisticated online banking fraud on extremely imbalanced data. *World Wide Web*, 16(4), 449–475.

[13] Whitrow, C., Hand, D.J., Juszczak, P., Weston, D., & Adams, N.M. (2009). Transaction aggregation as a strategy for credit card fraud detection. *Data Mining and Knowledge Discovery*, 18(1), 30–55.