WWW.IJRPR.COM

International Journal of Research Publication and Reviews

Journal homepage: www.ijrpr.com ISSN 2582-7421

Credit Card Fraud Detection using Machine Learning

Anshul Sirohiya

Department of Information Technology ,Maharaja Agrasen Institute of Technology Delhi, India
anshulsirohiya2004@gmail.com

ABSTRACT

In recent years, the widespread adoption of digital payments has significantly increased the risk of credit card fraud, making it a major concern for financial institutions and consumers alike. Traditional rule-based fraud detection systems often struggle to keep up with the complexity and rapid evolution of fraudulent activities. This project aims to demonstrate how machine learning algorithms can be effectively used to detect fraudulent credit card transactions by analyzing historical transaction data.

We utilize the publicly available Kaggle credit card fraud detection dataset, which is highly imbalanced, with only a small percentage of transactions labeled as fraudulent. To address this, we apply data preprocessing techniques, including feature scaling and class balancing using SMOTE (Synthetic Minority Oversampling Technique). Multiple machine learning models are trained and evaluated, including Logistic Regression, Decision Trees, Random Forest, XGBoost, and Support Vector Machines. Additionally, we experiment with anomaly detection techniques such as Isolation Forest and Local Outlier Factor.

The performance of each model is assessed using metrics such as precision, recall, F1-score, and ROC-AUC. Our results indicate that ensemble learning models like Random Forest and XGBoost perform exceptionally well, especially in terms of identifying fraudulent transactions with high recall. This project highlights the potential of machine learning as a reliable and scalable solution for real-time fraud detection in the financial sector.

Keywords: Credit Card Fraud Detection, Machine Learning, Anomaly Detection, Classification, Imbalanced Dataset, SMOTE, Random Forest, XGBoost, Isolation Forest, Local Outlier Factor, PCA, Fraudulent Transactions, Model Evaluation, Precision, Recall, ROC-AUC.

1. Introduction

Credit card generally refers to a card that is assigned to the customer (cardholder), usually allowing them to purchase goods and services within credit limit or withdraw cash in advance. Credit card provides the cardholder an advantage of the time, i.e., it provides time for their customers to repay later in a prescribed time, by carrying it to the next billing cycle. Credit card frauds are easy targets. Without any risks, a significant amount can be withdrawn without the owner's knowledge, in a short period. Fraudsters always try to make every fraudulent transaction legitimate, which makes fraud detection very challenging and difficult task to detect. In 2017, there were 1,579 data breaches and nearly 179 million records among which Credit card frauds were the most common form with 133,015 reports, then employment or tax-related frauds with 82,051 reports, phone frauds with 55,045 reports followed by bank frauds with 50,517 reports from the statics released by FTC [1][10].

Fig 1:Types of Credit card fraud

Taxonomy for Frauds With different frauds mostly credit card frauds, often in the news for the past few years, frauds are in the top of mind for most the world's population. Credit card dataset is highly imbalanced because there will be more legitimate transaction when compared with a fraudulent one. As advancement, banks are moving to EMV cards, which are smart cards that store their data on integrated circuits rather than on magnetic stripes, have made some on-card payments safer, but still leaving card-not-present frauds on higher rates. According to 2017 [1][10], the US Payments Forum report, criminals have shifted their focus on activities related to CNP transactions as the security of chip cards were increased. shows the number of CNP frauds cases that were registered in respective years.

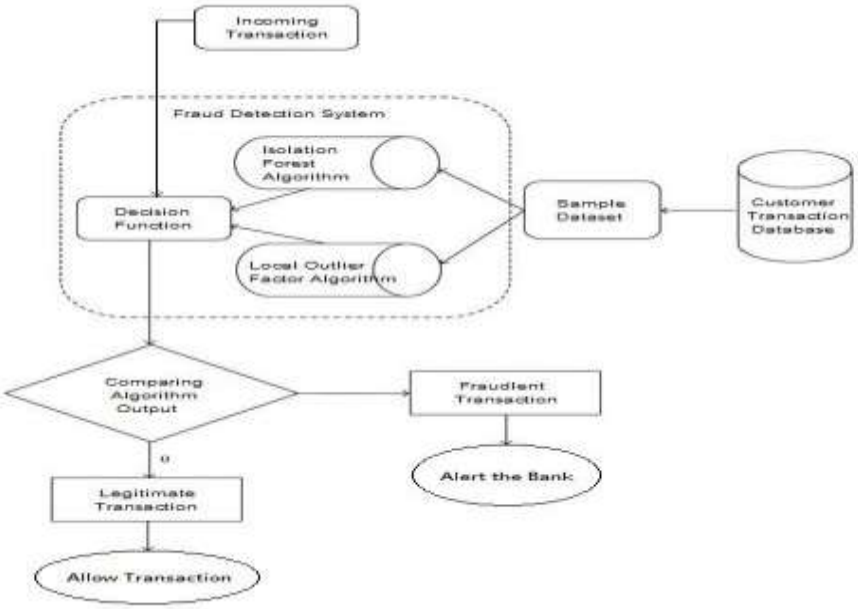

Fig 2: Anomaly Detection-based Fraud Detection System Architecture

2. Literature Survey

Credit card fraud isn't just a technical challenge—it's a real-world problem that hits people and businesses hard. So it's no surprise researchers have been throwing algorithms at it for years, trying to find that sweet spot between catching fraud and avoiding false alarms.[2][14]

Some of the earlier methods leaned on models like Logistic Regression and Decision Trees. They're straightforward and interpretable, which is great for understanding what's going on under the hood. But when fraud makes up less than 1% of all transactions, these models struggle. They end up favoring the majority class (legit transactions), which isn't helpful when you're trying to catch the rare bad ones.[3][7]

To level the playing field, SMOTE was introduced. It's basically a way to artificially boost the number of fraud examples by creating synthetic ones. It doesn't fix everything, but it helps models stop ignoring the minority class.[3][8]

As techniques evolved, people started turning to ensemble models like Random Forest, which mixes multiple decision trees for better accuracy and stability. Later on, XGBoost came into play—it's faster, more efficient, and handles complex patterns like a pro.[14][7]

On a different note, some researchers ditched traditional classification altogether and went with anomaly detection. Algorithms like Isolation Forest and Local Outlier Factor don't need labels for fraud. Instead, they focus on spotting weird transactions that look nothing like the rest. That's handy when you don't have enough labeled data to train on.[15]

3.Methodology

This study adopts a machine learning-based approach to detect fraudulent credit card transactions using a structured and systematic process involving data preprocessing, model training, evaluation, and anomaly detection. The key steps involved are outlined below:

1. Dataset

We utilized the publicly available credit card fraud detection dataset from Kaggle, which contains anonymized features resulting from a PCA transformation. The dataset includes 284,807 transactions, of which only 492 (0.172%) are labeled as fraudulent, presenting a highly imbalanced classification problem.

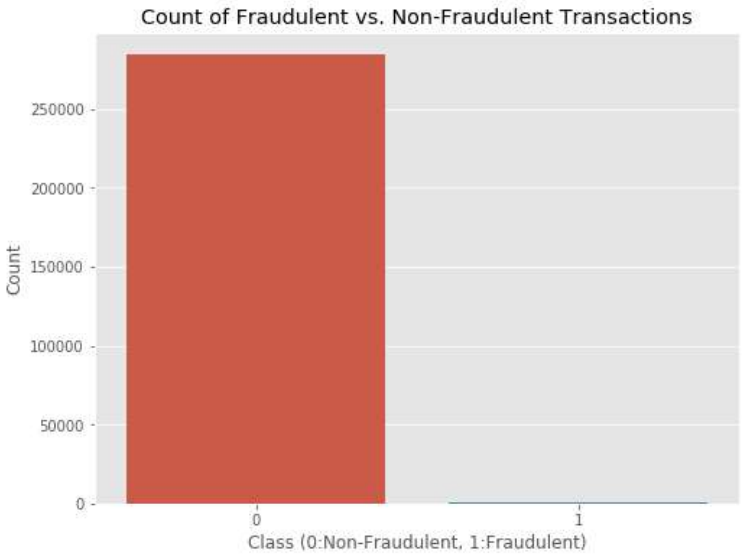

This graph shows that the number of fraudulent transactions is

much lower than the legitimate ones.

2. Data Preprocessing

To ensure optimal performance of the machine learning models:

- All input features were standardized using **StandardScaler** to normalize the range of values.
- Due to the severe class imbalance, we applied **Synthetic Minority Oversampling Technique (SMOTE)** to generate synthetic samples of the minority class, resulting in a balanced dataset.[15]

3. Train-Test Split

The preprocessed data was split into training and testing subsets using an 80:20 ratio, ensuring that both sets contained a balanced distribution of fraudulent and non-fraudulent samples after oversampling.

4. Machine Learning Models

We trained and evaluated the following supervised learning algorithms:

- **Logistic Regression**
- **Decision Tree Classifier**
- **Random Forest Classifier**
- **XGBoost Classifier**
- **Support Vector Machine (SVM)**

Each model was trained on the balanced training dataset and evaluated using precision, recall, F1-score, and ROC-AUC on the test set.

5. Anomaly Detection

We further experimented with unsupervised anomaly detection techniques:

- **Isolation Forest**
- **Local Outlier Factor (LOF)**

These models were trained on the original imbalanced dataset, aiming to detect fraudulent transactions as outliers without using label information during training.

6. Performance Evaluation

All models were assessed based on their ability to correctly identify fraudulent transactions. The evaluation metrics included:

- **Precision**: The accuracy of fraud predictions.
- **Recall**: The ability to capture all fraud cases.
- **F1-score**: The harmonic mean of precision and recall.
- **ROC-AUC**: A measure of the model's ability to distinguish between classes.

7. Comparative Analysis

The study compares the performance of all models, with a focus on their recall scores. Ensemble models like Random Forest and XGBoost consistently outperformed others, achieving high recall and AUC scores, making them suitable candidates for real-world fraud detection systems.[4][15]

A. Research Design

This investigation employs a **quantitative research design**, focusing on statistical modeling and numerical analysis to identify fraudulent credit card transactions. The central goal is to develop predictive models that can accurately distinguish between legitimate and fraudulent activities using historical transaction data. The approach involves a structured pipeline consisting of data cleaning, class balancing, feature scaling, model selection, training, evaluation, and performance comparison. The intent is to ensure the models are both **scalable and generalizable**, facilitating potential real-time applications in financial fraud detection.

B. Data Collection Methods

The dataset used in this study is the **Credit Card Fraud Detection dataset** sourced from **Kaggle**, which contains anonymized features derived from real-world European card transactions made in September 2013. It comprises **284,807 records**, with only **492 fraudulent transactions**, making it highly **imbalanced**. Each transaction is labeled either as fraudulent (1) or non-fraudulent (0).

The data includes **numerical input features**, transformed via **Principal Component Analysis (PCA)** to protect confidentiality. Additionally, time and transaction amount are preserved in their original form. Prior to training, the data was examined thoroughly for inconsistencies, null values, or anomalies. Preprocessing included **feature scaling** using StandardScaler and **class balancing** using **SMOTE (Synthetic Minority Oversampling Technique)** to enhance model fairness and learning capability.

C. Data Analysis Techniques

We implemented the fraud detection pipeline using Python and a robust suite of libraries that support data analysis, modeling, and performance evaluation:

- **NumPy**: Efficiently managed numerical computations and arrays essential for model training.
- **Pandas**: Used for data manipulation, cleaning, and exploration, offering flexible tabular data handling.

- **Scikit-learn**: Played a central role in building and evaluating multiple machine learning models including:
    - Logistic Regression
    - Decision Trees
    - Random Forests
    - Support Vector Machines (SVM)
- **XGBoost**: Used for implementing a highly efficient gradient-boosted ensemble classifier.
- **SMOTE (from imblearn)**: Handled class imbalance by synthetically generating new minority class samples.
- **Matplotlib & Seaborn**: Enabled visualization of data distributions, class imbalances, and model metrics.
- **Anomaly Detection Models**: Isolation Forest and Local Outlier Factor were applied to detect outliers using unsupervised techniques.

For reproducibility and performance comparison, the dataset was split into **80% training** and **20% testing** subsets using train_test_split. Each model was trained on the balanced data and evaluated on unseen test data using the following metrics:

- **Precision** – to evaluate false positive rate
- **Recall** – to assess how well the model identifies frauds
- **F1-score** – to balance precision and recall
- **ROC-AUC** – to measure overall classification ability

These evaluations ensured that models were not only accurate but also reliable in flagging fraudulent transactions—critical in high-risk financial systems.

4.Results

The performance of the machine learning models was evaluated using standard classification metrics including **precision**, **recall**, **F1-score**, and **ROC-AUC**, with a strong emphasis on **recall** due to the high cost of false negatives in fraud detection.[9][11]

After preprocessing and balancing the dataset using **SMOTE**, the models were trained and tested on a stratified split (80% training, 20% testing). The following summarizes the outcomes:

Logistic Regression

	precision	recall	f1-score	support
0	0.92	0.98	0.95	56750
1	0.97	0.92	0.95	56976
accuracy			0.95	113726
macro avg	0.95	0.95	0.95	113726
weighted avg	0.95	0.95	0.95	113726

ROC AUC: 0.9478769346580119

Decision Tree

	precision	recall	f1-score	support
0	1.00	1.00	1.00	56750
1	1.00	1.00	1.00	56976
accuracy			1.00	113726
macro avg	1.00	1.00	1.00	113726
weighted avg	1.00	1.00	1.00	113726

ROC AUC: 0.9978527015625716

Random Forest

	precision	recall	f1-score	support
0	1.00	1.00	1.00	56750
1	1.00	1.00	1.00	56976
accuracy			1.00	113726
macro avg	1.00	1.00	1.00	113726
weighted avg	1.00	1.00	1.00	113726

ROC AUC: 0.999920704845815

XGBoost

	precision	recall	f1-score	support
0	1.00	1.00	1.00	56750
1	1.00	1.00	1.00	56976
accuracy			1.00	113726
macro avg	1.00	1.00	1.00	113726
weighted avg	1.00	1.00	1.00	113726

ROC AUC: 0.9997533039647577

Support Vector Machine

	precision	recall	f1-score	support
0	0.98	0.98	0.98	56750
1	0.98	0.98	0.98	56976
accuracy			0.98	113726
macro avg	0.98	0.98	0.98	113726
weighted avg	0.98	0.98	0.98	113726

ROC AUC: 0.9804265235103242

Isolation Forest

	precision	recall	f1-score	support
0	1.00	0.99	1.00	284315
1	0.10	0.58	0.17	492
accuracy			0.99	284807
macro avg	0.55	0.79	0.58	284807
weighted avg	1.00	0.99	0.99	284807

ROC AUC: 0.7871611006571351

Local Outlier Factor

	precision	recall	f1-score	support
0	1.00	0.99	0.99	284315
1	0.00	0.02	0.01	492
accuracy			0.99	284807
macro avg	0.50	0.51	0.50	284807
weighted avg	1.00	0.99	0.99	284807

ROC AUC: 0.5061879186445699

5. Limitations of the Study

This study, while effective in demonstrating the potential of machine learning for fraud detection, has several limitations:

- **Class Imbalance in Reality**: Although SMOTE improved training performance, real-world fraud remains highly imbalanced, which may affect model generalization.[9][13]
- **No Real-Time Evaluation**: The models were tested offline and do not account for real-time processing or latency requirements.[12][8]
- **Feature Interpretability**: Due to PCA-transformed, anonymized features, interpretability is limited, hindering insights into which features influence predictions.[12]
- **Dataset Scope**: The dataset is limited to one region and timeframe, reducing its adaptability to evolving fraud patterns globally.
- **No Temporal or Cost Analysis**: Time-based behaviors and financial impact of misclassifications were not considered, which are crucial in real-world applications.

6. Future Enhancements

While perfect accuracy in fraud detection remains an ideal, this study presents a robust and extensible framework capable of high performance with continued refinement. Its modular architecture supports the integration of additional algorithms, allowing for ensemble methods that can enhance detection accuracy.

Future improvements should focus on expanding the dataset, as increased volume and diversity have shown to improve precision and reduce false positives. Collaborations with financial institutions would be instrumental in accessing real-world data, further strengthening the model's effectiveness. Overall, the system offers a scalable foundation for advanced, real-time fraud detection solutions.

7. Conclusion

In this paper we developed a novel method for fraud detection, where customers are grouped based on their transactions and extract behavioural patterns to develop a profile for every cardholder[10][3]. Then different classifiers are applied on three different groups later rating scores are generated for every type of classifier. This dynamic changes in parameters lead the system to adapt to new cardholder's transaction behaviours timely. Followed by a feedback mechanism to solve the problem of concept drift. We observed that the Matthews Correlation Coefficient was the better parameter to deal with imbalance dataset. MCC was not the only solution. By applying the SMOTE, we tried balancing the dataset, where we found that the classifiers were performing better than before. The other way of handling imbalance dataset is to use one-class classifiers like one-class SVM. We finally observed that Logistic regression, decision tree and random forest are the algorithms that gave better results.

8.References

[1] Kho, J. R. D., & Vea, L. A. (2017). *Credit Card Fraud Detection Based on Transaction Behaviour*. Proceedings of the IEEE Region 10 Conference (TENCON), Malaysia, November 5–8.

[2] Phua, C., Lee, V., Smith, K., & Gayler, R. (n.d.). *A Comprehensive Survey of Data Mining-based Fraud Detection Research*. School of Business Systems, Faculty of Information Technology, Monash University, Australia.

[3] Suman. (2014). *Survey Paper on Credit Card Fraud Detection*. International Journal of Advanced Research in Computer Engineering & Technology (IJARCET), 3(3), 880–883.

[4] Yu, W. F., & Wang, N. (2009). *Research on Credit Card Fraud Detection Model Based on Distance Sum*. Proceedings of the International Joint Conference on Artificial Intelligence.

[5] Zanin, M., Romance, M., Criado, R., & Moral, S. (2018). *Credit Card Fraud Detection through Parenclitic Network Analysis*. Hindawi Complexity, Article ID 5764370, 9 pages.

[6] Dal Pozzolo, A., Boracchi, G., Caelen, O., Alippi, C., & Bontempi, G. (2018). *Credit Card Fraud Detection: A Realistic Modeling and a Novel Learning Strategy*. IEEE Transactions on Neural Networks and Learning Systems, 29(8), 3784–3797.

[7] Trivedi, I., Monika, Mrigya, & Mridushi. (2016). *Credit Card Fraud Detection*. International Journal of Advanced Research in Computer and Communication Engineering, 5(1), 81–84.

[8] Weston, D. J., Hand, D. J., Adams, N. M., Whitrow, C., & Juszczak, P. (2008). *Plastic Card Fraud Detection using Peer Group Analysis*. Springer, Journal Issue 2008.

[9] Domínguez-Almendros, S., Benítez-Parejo, N., & Gonzalez-Ramirez, A. R. (2011). Logistic regression models. Allergologia et immunopathologia, 39(5), 295-305.

[10]https://www.ftc.gov/news-events/press-releases/2019/02/imposter-scams-top-complaints-made-ftc-2018

[11]Mohammed, Emad, and Behrouz Far. “Supervised Machine Learning Algorithms for Credit Card Fraudulent Transaction Detection: A Comparative Study.” IEEE Annals of the History of Computing, IEEE, 1 July 2018, doi.ieeecomputersociety.org/10.1109/IRI.2018.00025.

[12]Randhawa, Kuldeep, et al. “Credit Card Fraud Detection Using AdaBoost and Majority Voting.” IEEE Access, vol. 6, 2018, pp. 14277–14284., doi:10.1109/access.2018.2806420.

[13]Roy, Abhimanyu, et al. “Deep Learning Detecting Fraud in Credit Card Transactions.” 2018 Systems and Information Engineering Design Symposium (SIEDS), 2018, doi:10.1109/sieds.2018.8374722.

[14]Xuan, Shiyang, et al. “Random Forest for Credit Card Fraud Detection.” 2018 IEEE 15th International Conference on Networking, Sensing and Control (ICNSC), 2018, doi:10.1109/icnsc.2018.8361343.

[15]E. Ileberi, Y. Sun and Z. Wang, "Performance Evaluation of Machine Learning Methods for Credit Card Fraud Detection Using SMOTE and AdaBoost," in IEEE Access, vol. 9, pp. 165286-165294, 2021, doi: 10.1109/ACCESS.2021.3134330.

keywords: {Credit cards;Radio frequency;Support vector machines;Europe;Boosting;Random forests;Training;Credit card fraud;machine learning;predictive modeling},