International Journal for Multidisciplinary Research (IJFMR)
E-ISSN: 2582-2160 ● Website: www.ijfmr.com ● Email: editor@ijfmr.com

Machine Learning Approaches for Effective Credit Card Fraud Detection: Addressing Imbalance and Enhancing Accuracy

Shraddha Titare[1], Bhavitavya Isukapati[2], Diksha Waghmare[3], Snehal Jondhale[4], Suhas G. Salve[5]

[1,2,3,4]B.Tech Student, Computer Science Department, MGM's College of Engineering.
[5]Guide, Asst. Prof. (MTech. B.E.), Dept. of Computer Science & Engg., MGM's College of Engineering.

Abstract
Credit card fraud has emerged as a significant threat to the financial sector, driven by the rapid growth in online transactions and the evolving sophistication of fraudulent activities. This research aims to design and implement a machine learning-based solution capable of detecting fraudulent credit card transactions effectively. By addressing challenges such as dataset imbalance and false positives, the research employs preprocessing techniques including Synthetic Minority Oversampling Technique (SMOTE), along with advanced machine learning algorithms like Logistic Regression, XGBoost, and Isolation Forest. It highlights the potential of these models to enhance fraud detection accuracy and scalability, providing a practical and deployable tool for real-world applications. This comprehensive approach ensures the system is robust, adaptive, and user-friendly, paving the way for improved financial security and trust in digital payment systems.

Keywords: Credit Card Fraud Detection, Machine Learning Algorithms, Data Imbalance

1. Introduction

Credit card fraud detection is critical in ensuring the security of financial systems. The growing reliance on digital payments has transformed how financial transactions occur, but it has also opened new avenues for fraudsters to exploit vulnerabilities. Fraudulent activities have become increasingly sophisticated, ranging from phishing scams and identity theft to unauthorized transactions. These fraudulent activities lead to significant financial losses for institutions and individuals, as well as a decline in customer trust and confidence in digital payment systems. The complexity of detecting fraudulent transactions lies in the dynamic nature of fraud, where tactics evolve rapidly, and the vast majority of transactions are legitimate. This makes it challenging to identify anomalies accurately without disrupting genuine customer experiences. Additionally, the inherent imbalance in transaction datasets—where fraudulent activities are significantly outnumbered by legitimate ones—further complicates the task, as traditional machine learning models tend to favour the majority class.

This research focuses on designing a robust and adaptive machine learning system to detect fraudulent transactions in real-time. By leveraging advanced techniques for handling data imbalance, such as SMOTE, and employing state-of-the-art algorithms like Logistic Regression, XGBoost, and Isolation

Forest, the study aims to address the key challenges in fraud detection. The ultimate goal is to create a scalable and deployable solution that enhances the security of financial systems, minimizes financial losses, and restores user trust in the reliability of digital transactions.

2. Literature Review

Fraud detection has been a subject of extensive research, with approaches ranging from statistical methods to advanced machine learning algorithms. Logistic Regression and tree-based methods like XGBoost have shown promise in binary classification tasks. Unsupervised algorithms like Isolation Forest are effective in detecting anomalies in high-dimensional datasets. The domain of credit card fraud detection has significantly evolved with the advent of machine learning techniques. This literature review encapsulates approaches and their efficacy in combating fraudulent transactions.

2.1 Logistic Regression

Logistic regression is frequently employed for its simplicity and ability to provide probabilistic outcomes, making it a standard for binary classification tasks. Logistic Regression is a simple yet powerful machine learning algorithm widely used for binary classification problems, such as detecting fraudulent transactions in credit card data. It predicts the likelihood of a transaction being fraudulent by modelling the relationship between the target variable and the input features. The model outputs probabilities, which are then converted into class labels (fraudulent or non-fraudulent) based on a threshold.

The logistic function is defined as:

$P\ (Y{=}1) = 1 + e - z1$ (1)

Where $z = \beta 0 + \beta 1x1 + \beta 2x2 + \cdots + \beta nxn, and\ \beta$ represents the model coefficients.

The decision rule is:

- If $P(Y{=}1) \geq 0.5$: $Predict$ **1** $(Fraud)$.
- If $P(Y{=}1) < 0.5$: $Predict$ **0** $(Non{-}Fraud)$.

2.2 XGB Classifier

The XGBoost Classifier is an advanced and highly efficient machine learning algorithm widely used for classification tasks, including credit card fraud detection. XGBoost, which stands for Extreme Gradient Boosting, is based on the gradient boosting framework and builds an ensemble of decision trees sequentially to improve prediction accuracy. The algorithm performs exceptionally well with imbalanced datasets, a common characteristic of fraud detection, by emphasizing misclassified samples during training. Additionally, it provides feature importance scores, which help in understanding which factors contribute most to predicting fraud. These advantages make XGBoost a powerful and reliable tool for detecting fraudulent transactions in large, complex datasets.

2.3 Isolation Forest

Isolation Forest is an unsupervised anomaly detection algorithm that identifies anomalies by isolating data points. It works on the principle that anomalies are easier to isolate because they are few and different. The algorithm creates decision trees to isolate observations, measuring the number of splits required to separate a data point from the rest. Anomalous points typically require fewer splits, making them easier to detect. This method is efficient, scalable, and well-suited for high-dimensional datasets, offering robust performance in various anomaly detection scenarios.

The Isolation Forest implementation focuses on identifying anomalies (fraudulent transactions) in the dataset by leveraging an unsupervised learning technique. The process begins with the initialization of

the Isolation Forest model, where a contamination parameter is set. The contamination parameter (contamination=0.001) indicates the expected proportion of anomalies in the dataset. This is particularly useful in credit card fraud detection as fraudulent transactions typically form a very small portion of the overall data.

3. Methodology

The methodology of this research outlines a systematic approach to detecting fraudulent credit card transactions, integrating data analysis techniques and machine learning methods. A thorough understanding of the dataset and its characteristics serves as the foundation for building a robust fraud detection system. Key steps include feature engineering, and handling class imbalance, followed by the application of advanced algorithms. This approach ensures the model's reliability and scalability while maintaining a balance between accuracy and computational efficiency.

3.1 Dataset Overview

The dataset contains transactions made by credit cards in September 2013 by European cardholders. This dataset presents transactions that occurred in two days, where we have 492 frauds out of 284,807 transactions. The dataset is highly unbalanced, the positive class (frauds) account for 0.172% of all transactions.

The dataset used in this research is an extensive collection of credit card transaction records, comprising 284,807 entries and 31 features. The features include both original transaction attributes and those derived through advanced processing techniques such as Principal Component Analysis (PCA). While the anonymized features (V1 to V28) ensure privacy by obfuscating sensitive details, they retain their statistical significance, making them highly useful for machine learning applications. Alongside these transformed features, the dataset includes raw attributes such as Time, which denotes the elapsed time since the dataset's first transaction, and Amount, representing the financial value of the transaction. The target variable, Class, is a binary indicator where 0 signifies legitimate transactions and 1 flags fraudulent transactions.

3.2 Class Distribution

The extreme imbalance mirrors real-world scenarios where fraud is a relatively rare occurrence. While this rarity is a positive sign for financial systems in practice, it creates significant challenges for predictive modelling. During this phase, plotting the class distribution using bar charts or pie charts can vividly illustrate the degree of imbalance. This visualization serves as a diagnostic tool, highlighting the need for specialized strategies to handle the imbalance.

Figure 1: Class Distribution (Fraud vs Non-Fraud)

3.3 Feature Analysis

Effective feature analysis helps identify which features are most relevant for distinguishing between fraudulent and legitimate transactions, improving the overall performance and interpretability of machine learning models. Visualizations like histograms, density plots, and box plots are particularly useful for detecting outliers, skewness, and patterns in feature values. For example, in fraud detection, the Time feature, representing the time elapsed since the first transaction in the dataset, might reveal temporal patterns linked to fraudulent activity.

Figure 2: Fraud Transactions by Time

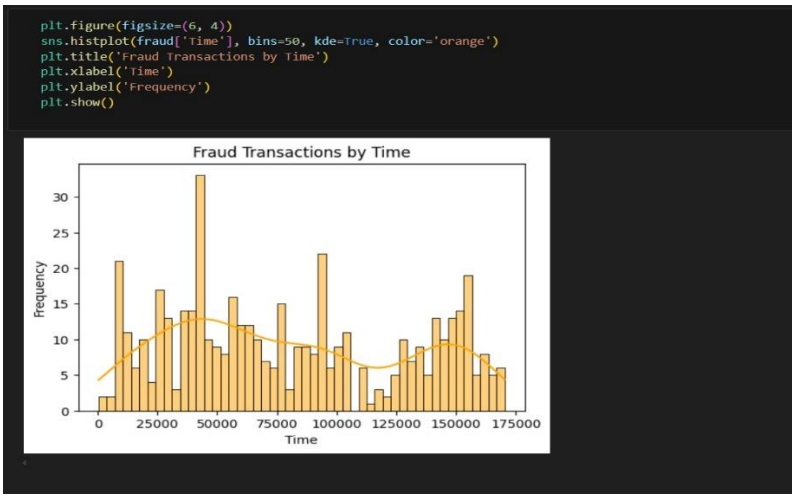

3.4 Correlation Heatmap

Correlation heatmaps are powerful visual tools used to explore and summarize the relationships between numerical features in a dataset. In fraud detection, where the dataset often includes numerous features, understanding these correlations is crucial for effective feature selection and model development. In fraud detection datasets, where features are sometimes pre-processed using dimensionality reduction techniques like Principal Component Analysis (PCA), correlation heatmaps can help assess the relationships among the PCA components or between the components and original features.

Figure 3: Correlation Heatmap

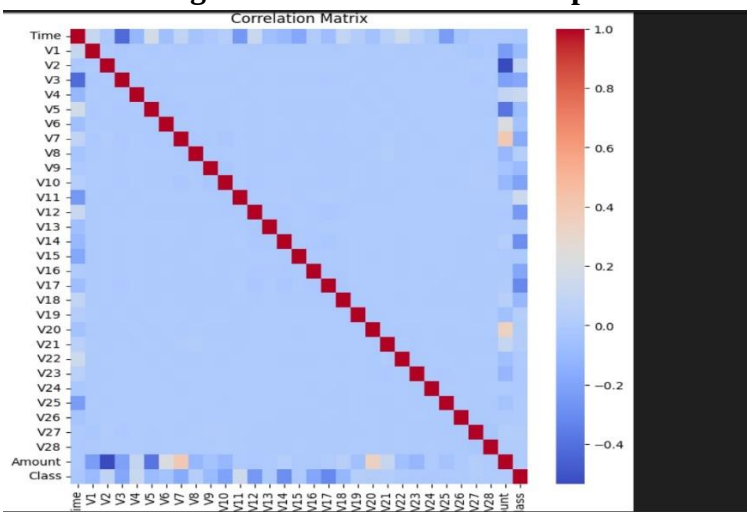

3.4 Outlier Detection

Outliers are data points that deviate substantially from the rest of the dataset, either due to errors, noise, or genuine anomalies. Identifying and handling these outliers is crucial because they can distort statistical analyses, skew machine learning models, and, in some cases, represent the very instances of interest, such as fraudulent transactions. For numerical features, visualization techniques such as box plots and scatter plots are particularly effective.

In fraud detection, outliers may represent high-value transactions, irregular purchasing patterns, or unusual combinations of features that deviate from typical behavior. For instance, an outlier in a feature like Amount could correspond to a transaction significantly larger than the average, which may warrant further investigation.

Figure 4: Outliers

3.5 Key Insights and Findings

The Classification Report and Confusion Matrix are essential tools for evaluating the performance of a credit card fraud detection model. The confusion matrix provides a detailed breakdown of the model's predictions by categorizing them into true positives, true negatives, false positives, and false negatives. This helps in visualizing how well the model distinguishes between fraudulent and legitimate transactions. The classification report complements the confusion matrix by summarizing key metrics such as precision, recall, and F1-score for each class.

Figure 5: Classification Report and Confusion Matrix

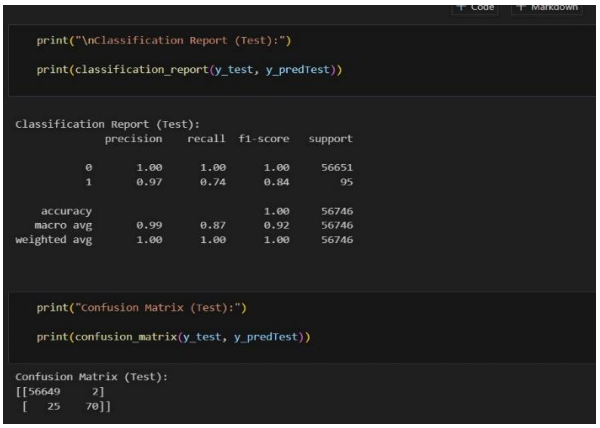

The Precision Score and Recall Score are pivotal metrics for understanding the effectiveness of a fraud detection model, particularly in imbalanced datasets. Precision measures the percentage of correctly identified fraudulent transactions out of all transactions predicted as fraudulent. Recall measures the percentage of actual fraudulent transactions correctly identified by the model.

Figure 6: Precision Score and Recall Score

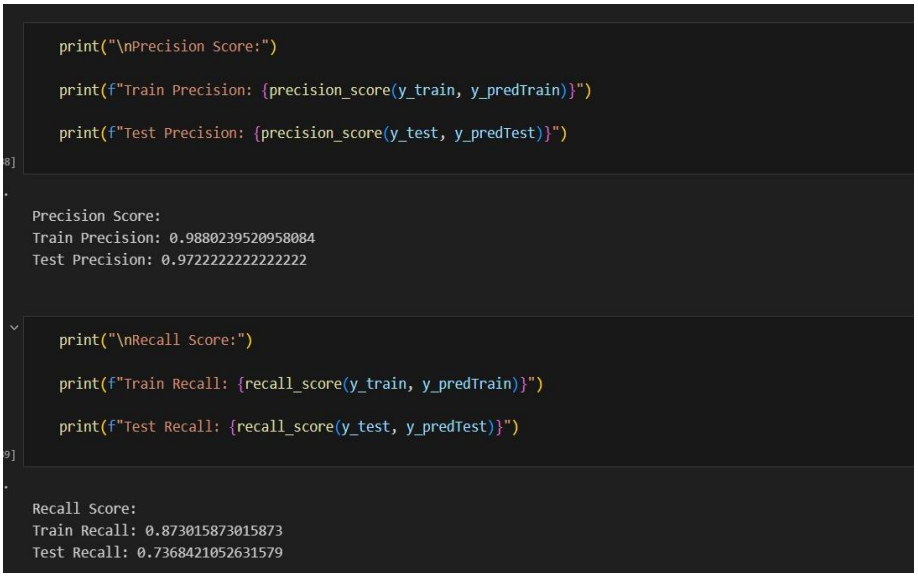

The Accuracy Score provides a clear and straightforward measure of a model's overall performance, representing the percentage of correct predictions (both fraudulent and legitimate) out of all predictions made. Accuracy serves as an important baseline metric to gauge how effectively the model identifies both types of transactions.

Fig 7: Accuracy Score

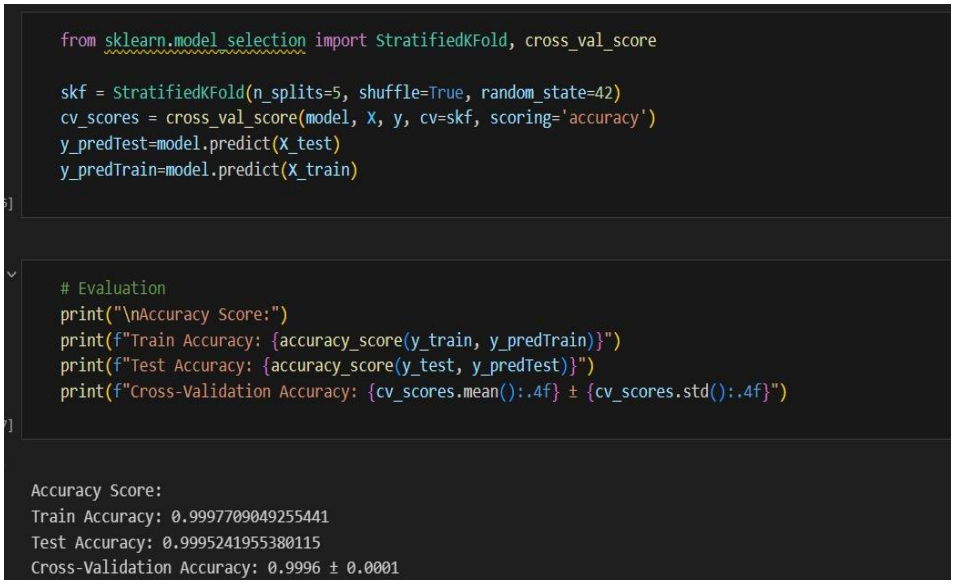

4. Results

By presenting tangible results and user interface examples, the outcome validates the success of the

machine learning model and its integration into a user-friendly application. It provides a comprehensive overview of how the system meets the needs outlined during the initial stages of the project and demonstrates its functionality. We addressed the seamless interaction between users and the system. The ability to upload datasets in a simple interface ensures accessibility, making the tool suitable for a wide audience, from technical experts to non-technical users. Additionally, the results table exemplifies how the system processes raw input data to deliver meaningful predictions, emphasizing the practicality and efficiency of the developed model. This foundation reflects the importance of translating complex computations into actionable insights, a critical component for fraud detection.

Figure 8: Input Data Forms

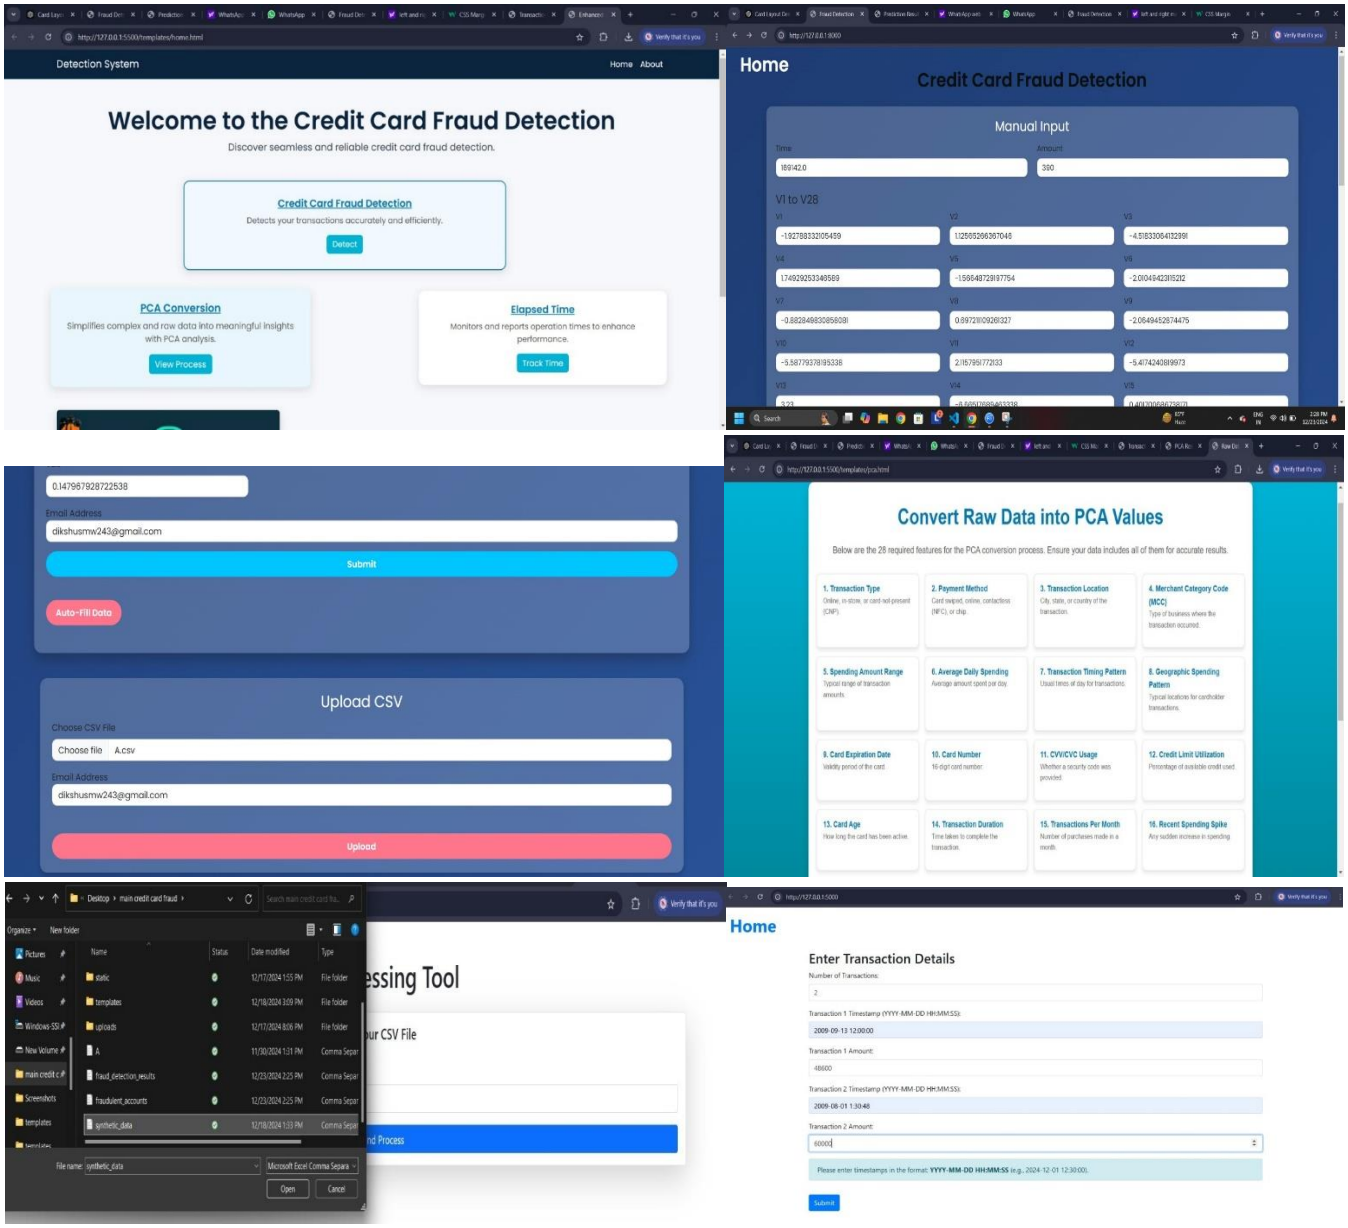

Figure 9: Prediction Results

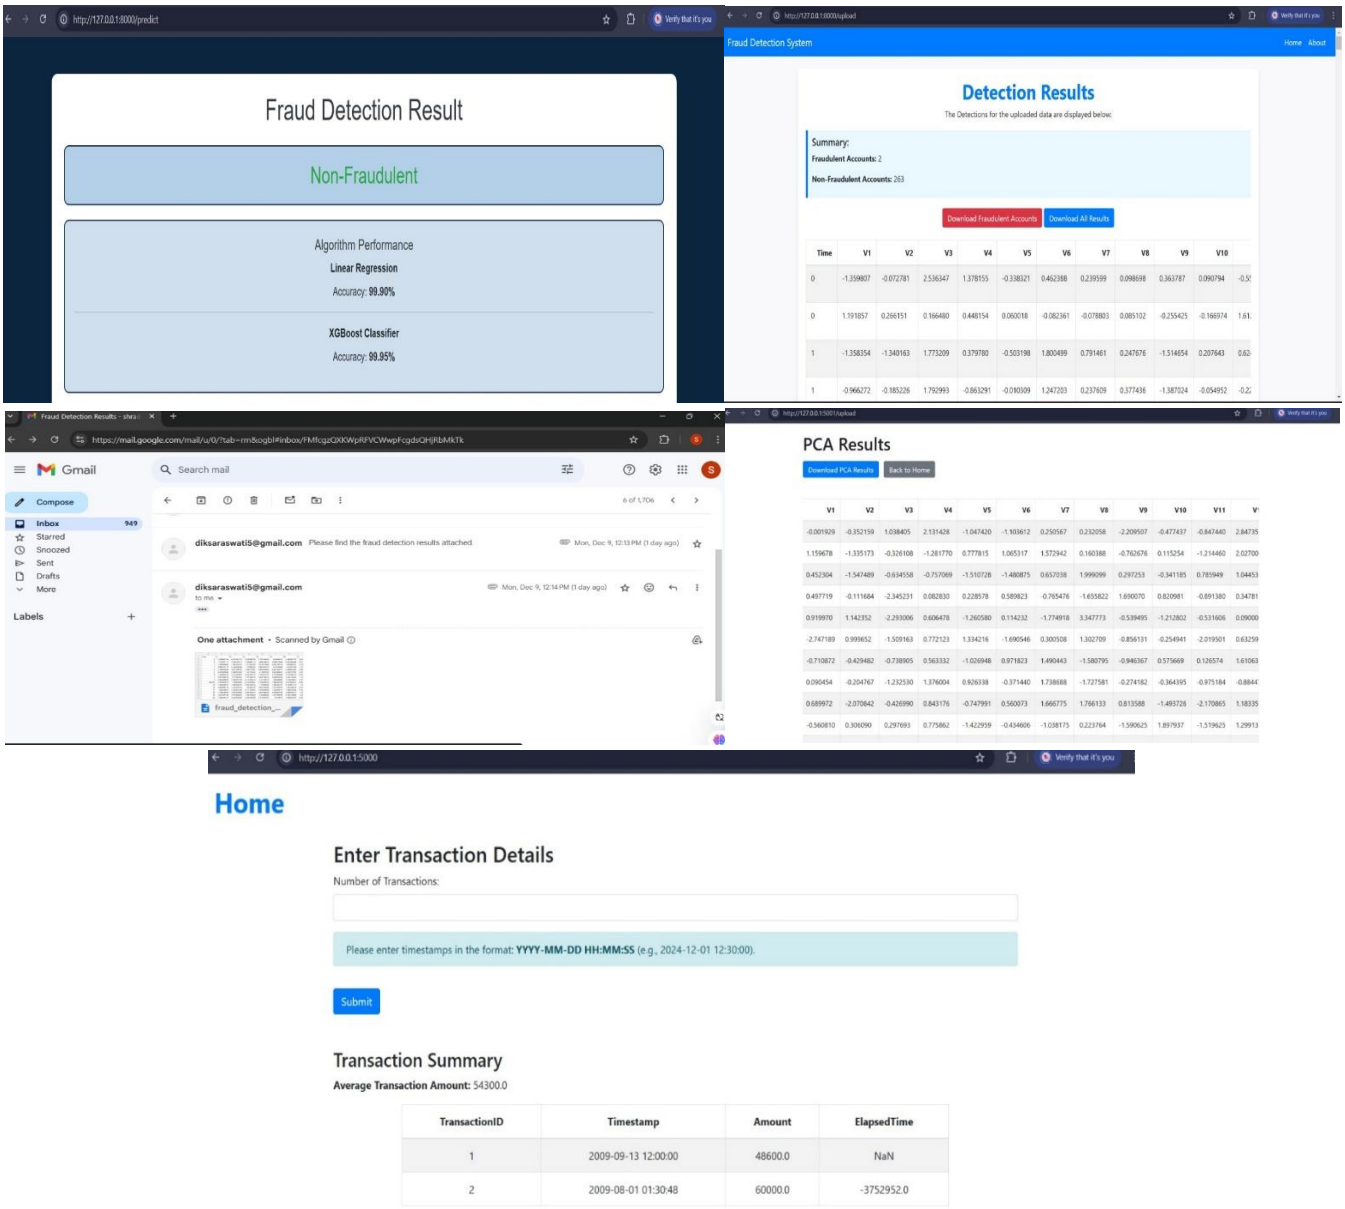

Conclusion

This research has effectively showcased the application of machine learning techniques to address the critical issue of credit card fraud detection. By tackling the inherent challenges of imbalanced datasets, the project emphasized the importance of preprocessing techniques, such as feature scaling, outlier detection, and balancing strategies like oversampling. These steps were instrumental in preparing the data and ensuring the model could effectively detect fraudulent transactions without being biased toward the majority class. The insights gained from exploratory data analysis and correlation studies further enhanced the understanding of the dataset, guiding feature selection and model optimization.

This research successfully demonstrated the potential of machine learning to tackle the complex problem of credit card fraud detection. It not only delivered a scalable and effective detection model but also highlighted the need for continuous learning and adaptation to keep pace with emerging threats. The

work provides a strong foundation for future innovations in fraud prevention, offering insights and techniques that can be extended to other domains of financial security.

References

1. A. Dal Pozzolo, G. Boracchi, O. Caelen, C. Alippi, and G. Bontempi, "Credit card fraud detection and concept-drift adaptation with delayed supervised information," in Proc. IEEE Int. Conf. on Data Mining Workshops (ICDMW), Singapore, 2018, pp. 938–947.
2. G. Lemaitre, F. Nogueira, and C. K. Aridas, "Imbalanced-learn: A Python toolbox to tackle the curse of imbalanced datasets in machine learning," Journal of Machine Learning Research, vol. 18, no. 1, pp. 559–563, Jan. 2017.
3. R. Patel and A. Thakur, "A study on credit card fraud detection using machine learning techniques," International Journal of Engineering Research & Technology (IJERT), vol. 8, no. 5, pp. 534–538, May 2019.
4. S. K. Sharma and M. P. Singh, "Credit card fraud detection using machine learning algorithms," in Proc. Int. Conf. on Computational Intelligence and Data Science (ICCIDS), Noida, India, 2018, pp. 1–5.