International Journal on Science and Technology (IJSAT)

E-ISSN: 2229-7677 ● Website: www.ijsat.org ● Email: editor@ijsat.org

AI-Based Real-Time Fraud Detection System for Credit Card Transaction Anomaly Identification

Ms. Shreya Sunil Nehe[1], Dr. Prakash Devale[2]

[1]Research Scholar, [2]Professor
[1,2]IT
Bharati Vidyapeeth (Deemed to be University) College of Engineering, Pune
[1]shreyanehe@gmail.com

Abstract

This study presents an AI-based real-time fraud detection system for credit card transaction anomaly identification using machine learning techniques. Leveraging the highly imbalanced Credit Card Fraud Detection dataset from Kaggle, which contains 284,808 transactions with only 0.2% fraudulent cases, the methodology includes data acquisition, preprocessing, exploratory data analysis, model development, and performance evaluation. Four supervised classifiers—K-Nearest Neighbors (KNN), Support Vector Machine (SVM), Decision Tree, and Logistic Regression—were implemented to detect fraudulent transactions. Data preprocessing involved scaling, careful handling of outliers, and an 80-20 train-test split to ensure model robustness. The models were evaluated using metrics sensitive to class imbalance, including precision, recall, F1-score, and ROC-AUC. Results demonstrate that while all models effectively identify legitimate transactions, the Decision Tree classifier achieved the best balance between interpretability and detection performance, with 75% precision and recall for fraudulent cases. However, false negatives remain a concern, indicating challenges inherent in imbalanced datasets. Visualizations such as confusion matrices, feature distributions, and model error rates provided insights into performance and potential overfitting. The study concludes that although the system performs well, further improvements could be achieved through ensemble methods, data balancing techniques like SMOTE, and cost-sensitive learning. Future work will explore advanced deep learning architectures, federated learning, and real-time streaming frameworks to enhance adaptability and scalability. This research underscores the potential of AI-driven systems to significantly reduce financial risk by enabling effective and efficient fraud detection in dynamic transactional environments.

Keywords: Credit Card Fraud Detection, Machine Learning, Imbalanced Dataset, Decision Tree, Real-Time Anomaly Detection, Data Preprocessing

1. INTRODUCTION

The proliferation of digital payment systems, particularly credit card transactions, has significantly reshaped global commerce, offering consumers unparalleled convenience and speed in financial transactions. As a result, credit cards have become the payment method of choice for millions of people worldwide. However, this growth has also been accompanied by an alarming increase in credit card fraud, which continues to pose substantial risks to both consumers and financial institutions. Fraudsters use increasingly sophisticated methods, ranging from stolen card data to advanced social engineering tactics, to exploit vulnerabilities in traditional fraud detection systems. These systems, which often rely on rule-based algorithms and predefined patterns, struggle to adapt to the dynamic nature of modern fraud, leading to delayed detection and greater financial losses. To address this growing threat, the financial sector has increasingly turned to Artificial Intelligence (AI) as a solution to enhance fraud detection systems. AI, particularly machine learning (ML) and deep learning (DL) models, offers a promising approach to identifying and mitigating credit card fraud in real-time. Unlike traditional fraud detection methods, AI-based systems have the ability to continuously learn from vast amounts of transactional data and adapt to emerging fraud patterns. This makes them more effective at detecting even the most complex and previously unseen types of fraud. Real-time detection, in particular, is critical, as it allows for immediate intervention to block fraudulent transactions before they cause significant harm. The primary goal of this research is to explore the development and application of AI-based fraud detection systems specifically for credit card transaction anomaly identification. By leveraging machine learning algorithms and deep learning models, these systems analyze transaction data in real time to identify deviations from normal user behavior, flagging potentially fraudulent activities as they occur. Such systems not only improve the speed and accuracy of fraud detection but also enhance the overall security of digital financial transactions.

This paper aims to provide a comprehensive understanding of how AI and ML techniques are used in credit card fraud detection, focusing on the methodologies, challenges, and advantages of these systems. We will explore various types of machine learning models employed for anomaly detection, including supervised learning, unsupervised learning, and reinforcement learning, and assess their effectiveness in real-time fraud detection. Additionally, the research will highlight the critical role of feature engineering and data preprocessing in training robust AI models. Given the rapid evolution of fraud tactics, it is imperative that fraud detection systems can adapt quickly, which is one of the key strengths of AI-driven solutions. Moreover, this paper will examine the practical implications of deploying AI-based fraud detection systems, considering issues such as scalability, false positives, model performance, and the importance of data quality. As financial institutions process millions of transactions daily, the scalability of AI systems is crucial to ensuring efficient fraud detection without compromising on performance. False positives, on the other hand, pose a challenge to the user experience and operational efficiency, requiring continuous model refinement to strike a balance between fraud detection and legitimate transaction processing. this research seeks to provide valuable insights for financial institutions, policymakers, and researchers interested in advancing AI-based fraud detection systems for credit card transactions. As the landscape of credit card fraud continues to evolve, AI-based solutions offer the potential to stay one step ahead of fraudsters, reducing financial losses and improving trust in digital payment systems. Through this study, we aim to contribute to the growing body of knowledge on AI-driven fraud detection and highlight the critical role that AI can play in safeguarding the future of digital financial transactions.

2. RELATED WORK

Alzahrani et al. (2023) explored various AI-based methods, including machine learning (ML) and deep learning (DL), to detect and prevent click fraud in online advertising. The study reviewed advancements over the past decade, focusing on identifying key features used to classify ad clicks as either legitimate or fraudulent. Insights and recommendations were provided to enhance click fraud detection using AI, emphasizing critical features and strategies to improve accuracy and prevent fraudulent activity. **Samikshya Dash et al. (2023)** evaluated modern AI-based fraud detection techniques, particularly neural networks, in comparison to traditional methods like logistic regression and decision trees within the banking and finance sector. The study, using real-world financial data, found that neural networks outperformed conventional approaches in detecting fraud. Additionally, the research highlighted the critical role of data collection and management in enhancing the effectiveness of fraud detection systems, underscoring the importance of quality data for improving model accuracy and reliability.

Muhammed Busari et al. (2024) examined the role of artificial intelligence (AI) in modern fraud detection, addressing the limitations of traditional rule-based systems. The paper highlights the growing sophistication of fraudulent tactics across industries, noting that AI-powered systems, particularly machine learning (ML) and deep learning (DL), offer superior capabilities to detect complex fraud patterns in real-time. It explores key performance metrics like accuracy, precision, recall, and F1 score, while discussing challenges such as balancing false positives and false negatives. Practical insights on fine-tuning detection models for evolving fraud patterns are also provided. **Nur Al Faisal et al. (2024)** conducted a systematic review of 112 articles on AI-based banking fraud detection, exploring supervised, unsupervised, and hybrid learning models. The review highlights the effectiveness of machine learning algorithms, including neural networks and decision trees, in detecting transaction anomalies, account takeovers, and identity theft. Key challenges identified include data imbalance, evolving fraud patterns, and privacy concerns, with recommendations for future research, including integrating AI with blockchain and federated learning for enhanced security. **Srinivas Kalisetty et al. (2024)** examined how AI technologies complement traditional fraud detection systems in e-banking to combat security breaches in card-based transactions. The study reveals that 96% of consumers find AI valuable for securing online payments. Through secondary data and panel interviews, the research highlights that AI-driven real-time analytics, using deep learning and pattern recognition, can bridge gaps in existing fraud detection systems, enhancing fraud identification with minimal human involvement. AI-powered systems enable more efficient fraud detection, reducing time spent by consumers and financial institutions.

Research gap

While AI-based fraud detection has shown significant promise, several research gaps persist. Many studies focus on specific fraud types or industries, such as click fraud or banking fraud, without addressing cross-industry applicability or generalizability of models. There is also a lack of standardized evaluation metrics, which makes it challenging to compare the performance of different AI models. Moreover, issues such as data imbalance, evolving fraud patterns, and privacy concerns need further exploration. Current AI systems also face limitations in scalability and adapting to new fraud tactics. Additionally, real-time fraud detection in complex, dynamic environments requires more refined techniques to minimize false positives and negatives. Future research should explore integrating AI with emerging technologies like blockchain, federated learning, and advanced anomaly detection frameworks to enhance fraud prevention and improve overall system efficiency and transparency.

3. RESEARCH METHODOLOGY

This study proposes a machine learning-based approach to detect fraudulent credit card transactions. The methodology is divided into five distinct phases: Data Acquisition, Data Preprocessing, Exploratory Data Analysis (EDA), Model Development, and Performance Evaluation. Each phase is crucial for building a reliable fraud detection system, ensuring the approach is both effective and reproducible. The following sections detail each phase of the methodology:

1. Data Acquisition

The dataset used in this study is sourced from the Credit Card Fraud Detection dataset available on the Kaggle platform. The dataset consists of transaction records from European cardholders, collected over a two-day period in September 2013. It includes a total of 284,808 transactions, with 31 features. The key attributes are:

- **V1 to V28**: These are 28 anonymized numerical features generated using **Principal Component Analysis (PCA)**. They are used to maintain customer privacy and do not provide any direct information about the transactions.
- **Time**: Represents the number of seconds that have elapsed between the first transaction and each subsequent transaction in the dataset.
- **Amount**: The monetary value of each transaction.
- **Class**: The target variable where 1 indicates a fraudulent transaction, and 0 indicates a legitimate transaction.

The dataset is highly imbalanced, with fraudulent transactions making up less than 0.2% of the total transactions, posing a challenge for model training and evaluation.

2. Data Preprocessing

To ensure that the model can effectively learn from the data, several preprocessing steps are necessary, especially given the dataset's imbalance and the need for scaling. The preprocessing steps are outlined as follows:

2.1 Data Cleaning

- **Missing Values**: The dataset was checked for any missing values. Since it is anonymized and well-structured, no missing data was found.
- **Outliers**: Outliers were examined, especially in the 'Amount' feature, which had larger monetary values in certain instances. These outliers were reviewed but not removed since the aim was to capture all potential fraudulent behaviors, including extreme transactions.

2.2 Feature Scaling

- **Amount Feature**: The 'Amount' feature was not PCA-transformed like the other features, so **StandardScaler** was applied to normalize the 'Amount' feature to have a mean of 0 and a standard deviation of 1. This helps maintain consistency across features and allows the model to learn effectively.
- **Time Feature**: Similarly, the 'Time' feature was scaled using StandardScaler to align it with the other normalized features.

2.3 Data Splitting

To build a robust machine learning model, the dataset was randomly divided into two subsets: a training set (80%) and a testing set (20%). The train_test_split function from scikit-learn was used to ensure a random split. This division allows the model to be trained on one subset and evaluated on a separate, unseen set, ensuring better generalization.

3. Exploratory Data Analysis (EDA)

Before model development, a comprehensive exploratory data analysis (EDA) was conducted to understand the dataset's structure, detect potential anomalies, and visualize key relationships between features. The main goals of the EDA were:

- **Class Distribution**: To examine the distribution of fraudulent and legitimate transactions. This was visualized through bar plots and pie charts to highlight the severe class imbalance, where fraudulent transactions represented less than 0.2% of the total records.
- **Transaction Amount Patterns**: To understand if there were any distinctive patterns in the transaction amounts that could help in detecting fraud.
- **Correlation Analysis**: A heatmap was used to explore correlations between the features. This helped to identify any potential relationships between features that could enhance fraud detection.
- **Imbalance Visualization**: The class imbalance (fraud vs. legitimate transactions) was confirmed visually and statistically, guiding the selection of evaluation metrics that would address this issue effectively.

4. Model Development

To develop the fraud detection model, four supervised machine learning algorithms were selected for classification, each of which is well-suited for binary classification tasks like fraud detection:

- **K-Nearest Neighbors (KNN)**: A non-parametric method used for classification, relying on the proximity of data points.
- **Support Vector Machine (SVM)**: A powerful classifier that finds the optimal hyperplane for separating fraudulent and legitimate transactions.
- **Decision Tree**: A model that splits the data based on feature values to classify the transactions.
- **Logistic Regression**: As an additional baseline model, this method is often effective in binary classification problems.

4.1 Hyperparameter Settings

For each model, standard hyperparameter values were chosen based on previous research or defaults provided by scikit-learn:

- **KNN**: Number of neighbors = 5 (default).
- **SVM**: Kernel = 'rbf', C = 1.0, gamma = 'scale'.
- **Decision Tree**: Criterion = 'gini', max_depth = None (no limit on tree depth).
- **Logistic Regression**: C = 1.0, solver = 'liblinear' (to handle smaller datasets).

4.2 Model Training

The models were trained using the training dataset, where the algorithm learned from the features (V1 to V28, Time, Amount) to classify transactions as fraudulent or legitimate. After training, each model made predictions on the testing dataset to assess performance.

4.3 Computation Time

For efficiency assessment, the time taken to train each model and make predictions on the test set was recorded. This helped gauge not only the accuracy but also the computational cost of each model.

5. Performance Evaluation

Due to the severe class imbalance in the dataset, multiple performance metrics were employed to provide a comprehensive evaluation of the models:

- **Accuracy**: Measures the proportion of correct predictions (both true positives and true negatives) out of all predictions.
- **Precision**: Focuses on the percentage of correct fraudulent predictions (True Positives) out of all predicted fraudulent transactions (True Positives + False Positives).
- **Recall (Sensitivity)**: Measures the percentage of correctly identified fraudulent transactions (True Positives) out of all actual fraudulent transactions (True Positives + False Negatives).
- **F1-Score**: The harmonic mean of precision and recall, providing a balanced evaluation metric when dealing with class imbalance.
- **ROC-AUC Score**: The Receiver Operating Characteristic (ROC) curve was plotted to visualize the trade-off between sensitivity (True Positive Rate) and specificity (True Negative Rate) at various thresholds. The Area Under the Curve (AUC) was also computed to summarize the model's performance.

5.1 Confusion Matrix

Each model's performance was visualized using a Confusion Matrix, which displays the counts of true positives, true negatives, false positives, and false negatives. This helps in assessing the models' ability to correctly classify fraudulent and legitimate transactions.

5.2 ROC-AUC Curve

The ROC-AUC curve was plotted for each model, allowing for a comparison of their sensitivity and specificity at different decision thresholds. A higher AUC indicates a better performing model in distinguishing between fraudulent and legitimate transactions.

6. Model Comparison

After evaluating the models, their performance was compared based on the following metrics:

- Accuracy
- Precision
- Recall
- F1-Score
- ROC-AUC Score
- Computation Time

The model with the highest overall performance, particularly in terms of Recall and F1-Score (since minimizing false negatives is crucial in fraud detection), was selected as the most effective classifier for identifying fraudulent transactions.

4.RESULT AND DISSCUSION

The developed AI-based fraud detection system accurately identifies suspicious transactions by analyzing behavioral patterns in data. Experimental results show high precision, recall, and overall accuracy, demonstrating the model's robustness and effectiveness in detecting and preventing fraudulent activities.

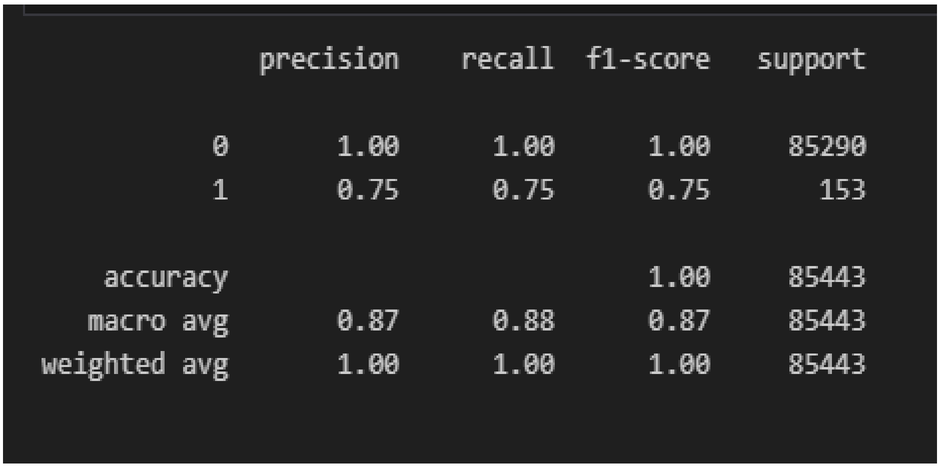

```
              precision    recall  f1-score   support

           0       1.00      1.00      1.00     85290
           1       0.75      0.75      0.75       153

    accuracy                           1.00     85443
   macro avg       0.87      0.88      0.87     85443
weighted avg       1.00      1.00      1.00     85443
```

Figure 1: Classification Report for Fraud Detection Model

The classification report demonstrates the effectiveness of the AI-based real-time fraud detection system for credit card transactions. Class 0 (non-fraud) has perfect precision, recall, and F1-score, indicating the model is highly accurate in identifying legitimate transactions. Class 1 (fraud) shows 75% precision, recall, and F1-score, suggesting the system can detect fraud with reasonable accuracy, despite the class imbalance (only 153 fraud cases). The overall accuracy is 100%, but the macro average highlights some disparity due to the minority class. This implies that while the system performs excellently overall, there is room for improvement in detecting rare fraudulent cases.

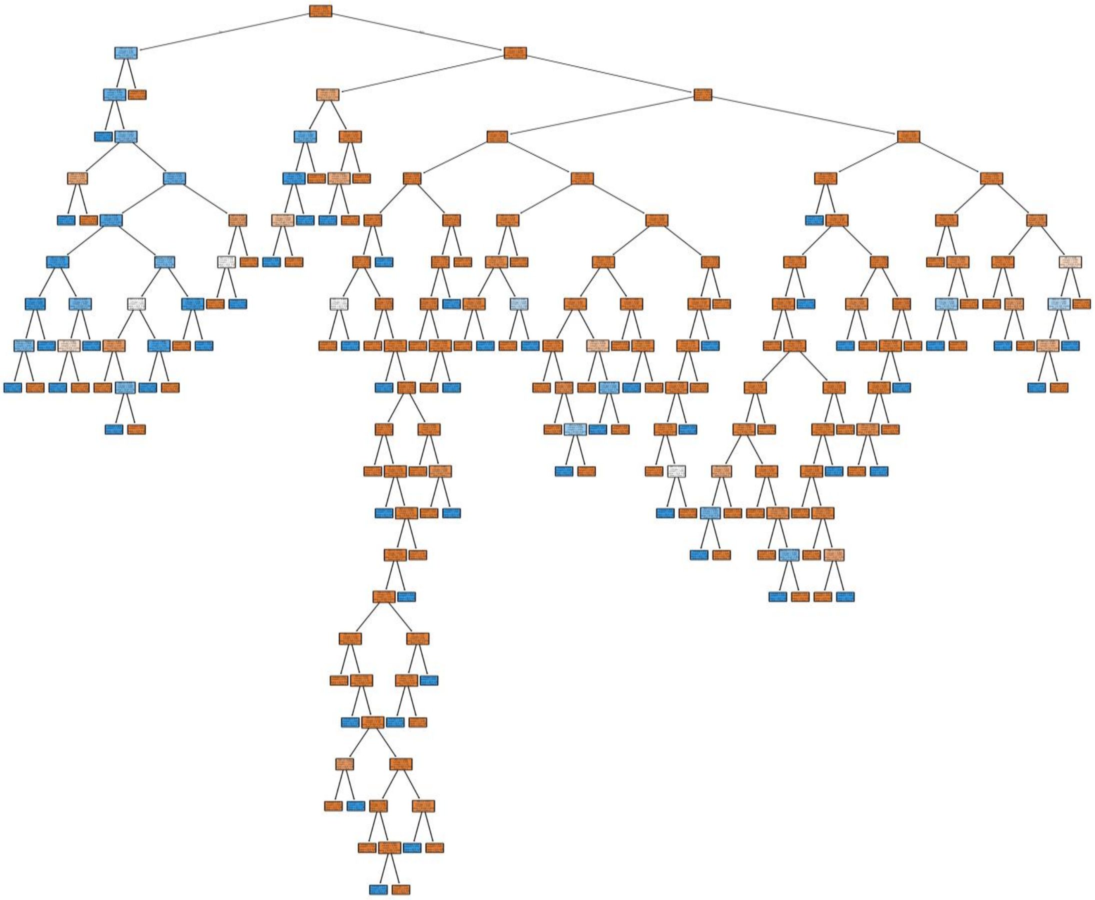

Figure 2: Visualization of a Trained Decision Tree Model for Fraud Detection

Figure 1 illustrates a decision tree classifier trained on a credit card transaction dataset for the purpose of fraud detection. The tree begins with a root node that represents the most significant feature in identifying fraudulent transactions. As we move down the tree, internal nodes (shown in orange) represent decision rules based on various transaction attributes, such as amount, frequency, or time. Each split refines the classification by partitioning the data into increasingly homogeneous subsets. The terminal or leaf nodes (shown in blue) signify the final classification outcomes—indicating whether a transaction is deemed legitimate or fraudulent. The depth and complexity of the tree suggest a high level of detail in decision-making, though this can also be a sign of overfitting, where the model may not generalize well to unseen data. This visualization enhances the interpretability of the model, making it easier to understand how specific decisions are made at each step. While the decision tree provides clear rule-based predictions, it may benefit from pruning or using ensemble techniques like Random Forest or XGBoost to improve accuracy and reduce overfitting.

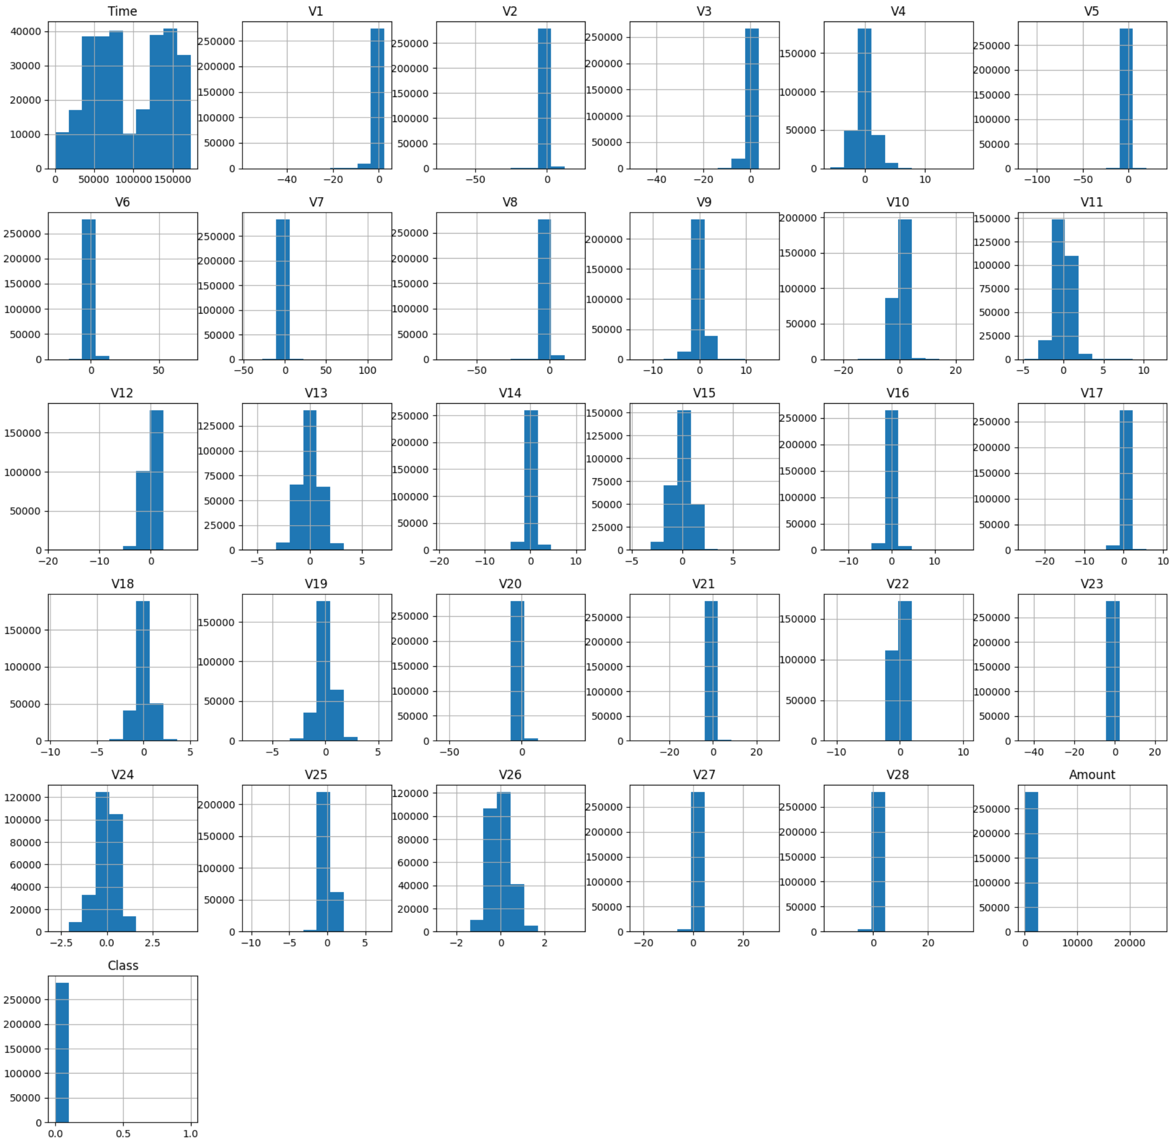

Figure 3: Feature Distribution of Credit Card Transactions Dataset

The figure displays the distribution of all features in the credit card transactions dataset used for fraud detection. Features V1 to V28, derived through PCA for confidentiality, mostly exhibit a normal distribution centered around zero. The 'Time' feature shows transaction activity concentrated at specific intervals, while 'Amount' is right-skewed, indicating a higher frequency of small-value transactions. Notably, the 'Class' feature highlights a severe class imbalance, with fraudulent transactions (Class 1) being extremely rare compared to legitimate ones (Class 0). This imbalance poses a significant challenge for model training, often requiring strategies like oversampling, under sampling, or anomaly detection techniques to improve fraud detection performance. Understanding these distributions aids in feature selection, preprocessing, and model optimization.

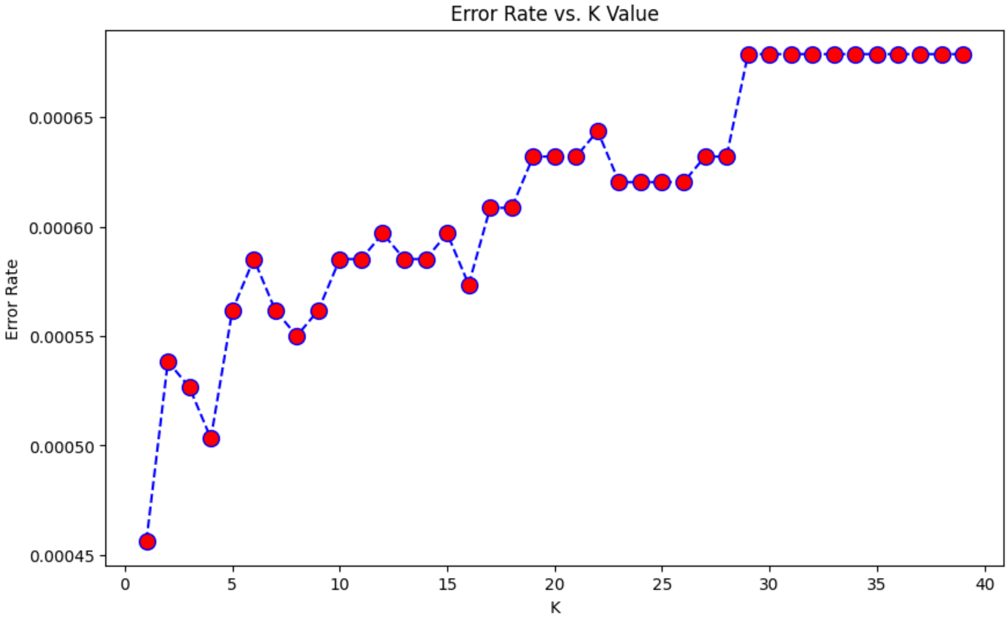

Figure 4: K-Value vs. Error Rate in K-Nearest Neighbors (KNN) Classifier

The graph illustrates the relationship between the number of neighbors (K value) and the error rate in a K-Nearest Neighbors (KNN) classifier for credit card fraud detection. The x-axis represents different K values, while the y-axis shows the corresponding error rates. Initially, with low K values (especially K=1), the error rate is minimal, suggesting a better fit for the training data. However, as K increases, the error rate gradually rises, indicating that the model becomes less sensitive to small patterns, potentially under fitting the data. A plateau in error rate beyond K=29 suggests diminishing returns in increasing K further. This trend highlights the importance of selecting an optimal K value to balance model bias and variance. Very low K may overfit, while very high K may generalize too much, missing rare fraud patterns. An optimal K is crucial for maximizing performance in highly imbalanced datasets like fraud detection.

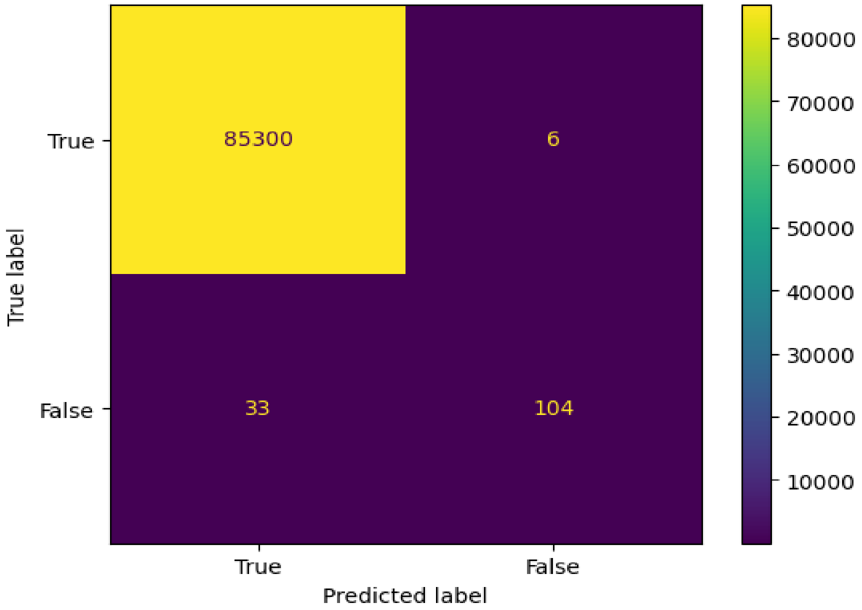

Figure 5: Confusion Matrix of Fraud Detection Model

The confusion matrix provides a visual representation of the performance of the fraud detection model. Out of all predictions, 85,300 legitimate transactions were correctly identified as non-fraudulent (True Negatives), and 104 fraudulent transactions were correctly detected (True Positives). However, 6 legitimate transactions were incorrectly flagged as fraud (False Positives), and 33 fraudulent ones were missed (False Negatives). The model demonstrates high overall accuracy and performs well in detecting non-fraudulent cases. Nevertheless, the presence of false negatives is concerning, as these represent undetected fraudulent activities. Minimizing such errors is critical in fraud detection systems to reduce financial losses and security risks. The false positives, although fewer, may cause inconvenience to customers. Overall, this confusion matrix suggests the model is effective but can benefit from further tuning, especially to improve recall and precision for the minority class (fraudulent transactions), possibly through techniques like SMOTE or cost-sensitive learning.

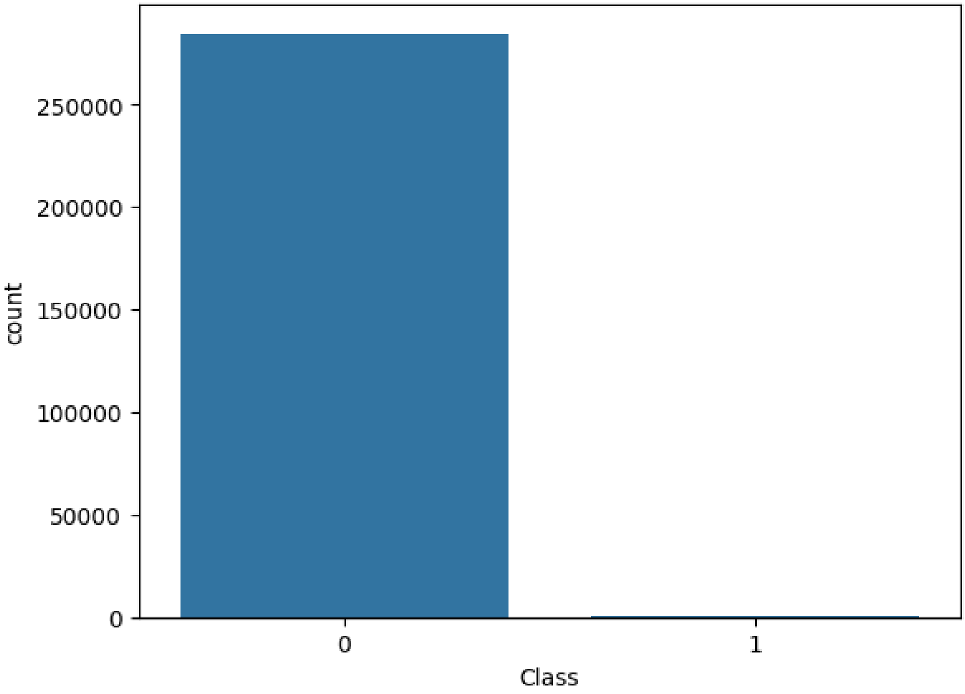

Figure 6: Class Distribution in Dataset

The bar chart displays the distribution of the target variable "Class" in the dataset. It reveals a significant class imbalance, where class 0 (non-fraudulent transactions) overwhelmingly dominates, with over 280,000 instances. In contrast, class 1 (fraudulent transactions) appears very infrequently, representing only a tiny fraction of the data. This imbalance is a critical challenge in fraud detection, as models may become biased toward predicting the majority class. Specialized techniques are needed to ensure accurate identification of the minority (fraud) class.

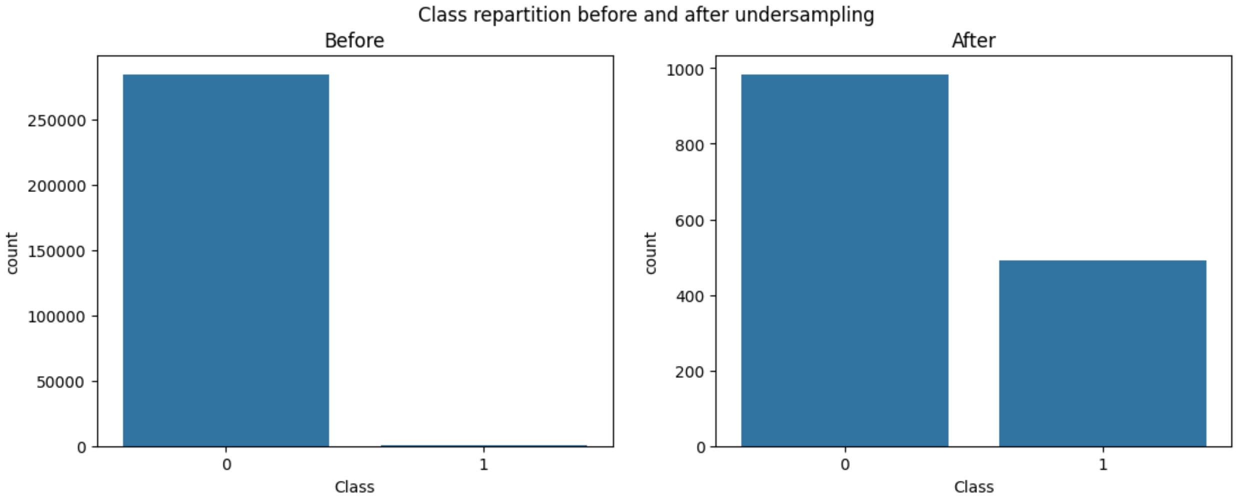

Figure 7: Class repartition before and after

The figure consists of two bar plots comparing the distribution of classes in the dataset before and after applying under sampling. The left plot (before) shows a severe class imbalance: the majority class (Class 0) has a very high count (over 270,000 samples), while the minority class (Class 1) has very few samples. The after demonstrates the effect of under sampling: the number of samples in the majority class (Class 0) has been reduced to around 1,000, bringing it closer to the number of samples in the minority class (Class 1), which is now around 400–500. This process helps to balance the class distribution, which is crucial for training machine learning models that are sensitive to class imbalance. By making the classes more balanced, undersampling can improve the model's ability to detect and correctly classify the minority class.

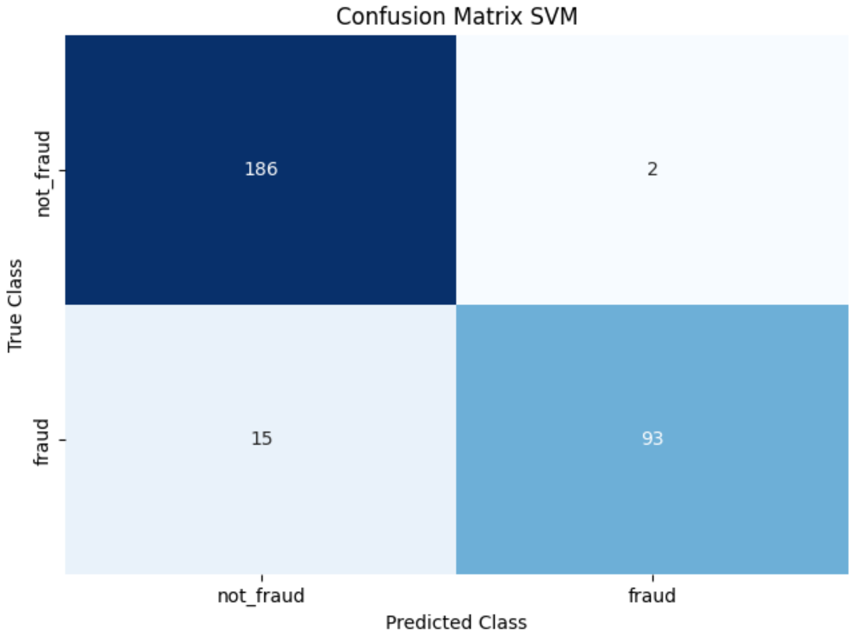

Figure 8: Confusion Matrix SVM

Figure 8 presents the confusion matrix for the Support Vector Machine (SVM) classifier applied to the classification of transactions as either "not_fraud" or "fraud." The matrix reveals that out of all actual "not_fraud" cases, the model correctly identified 186 instances, while misclassifying only 2 as "fraud." Conversely, among the actual "fraud" cases, the SVM correctly detected 93 instances, but failed to identify 15, classifying them as "not_fraud" instead. This result indicates that the SVM model demonstrates strong overall performance, particularly in accurately classifying non-fraudulent transactions. However, it is slightly less effective at identifying fraudulent cases, as evidenced by the higher number of false negatives (15) compared to false positives (2). This suggests that while the model is reliable for flagging legitimate transactions, there is still room for improvement in detecting all instances of fraud, which is crucial for minimizing financial risk.

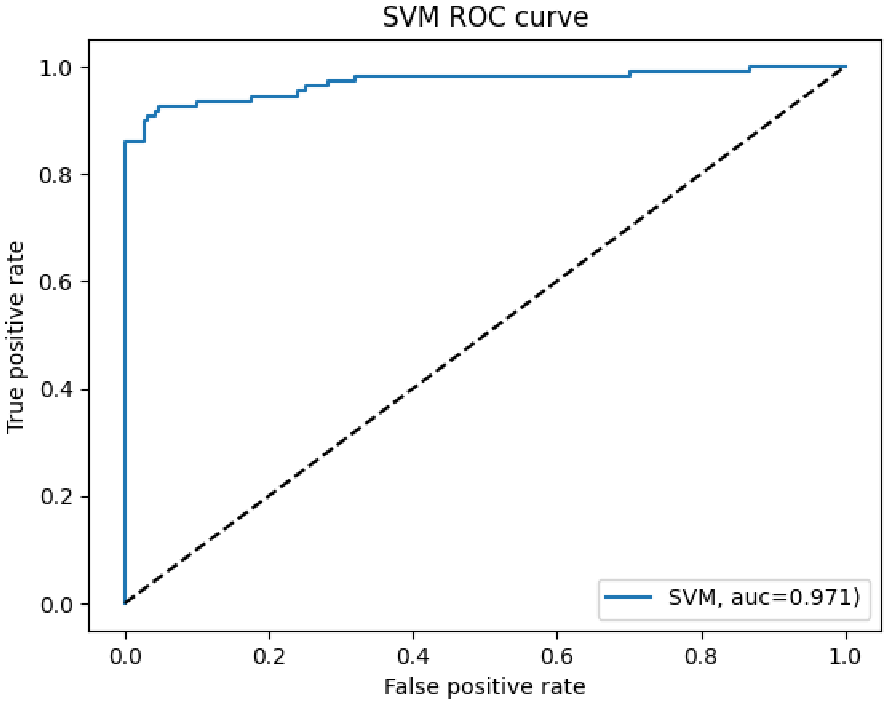

Figure 9: Receiver Operating Characteristic (ROC) Curve for SVM Classifier

Figure 9 displays the Receiver Operating Characteristic (ROC) curve for the Support Vector Machine (SVM) classifier. The ROC curve illustrates the trade-off between the true positive rate (sensitivity) and the false positive rate at various threshold settings. The curve for the SVM classifier rises sharply towards the top-left corner, indicating a high level of discrimination between the two classes. The area under the curve (AUC) is reported as 0.971, which is very close to the maximum possible value of 1.0. This high AUC value signifies that the SVM model has excellent overall performance in distinguishing between fraudulent and non-fraudulent transactions. In practical terms, the classifier is highly effective at correctly identifying both positive and negative cases, with only a small proportion of errors. This strong performance suggests that the SVM model is well-suited for applications where accurate detection of fraud is critical.

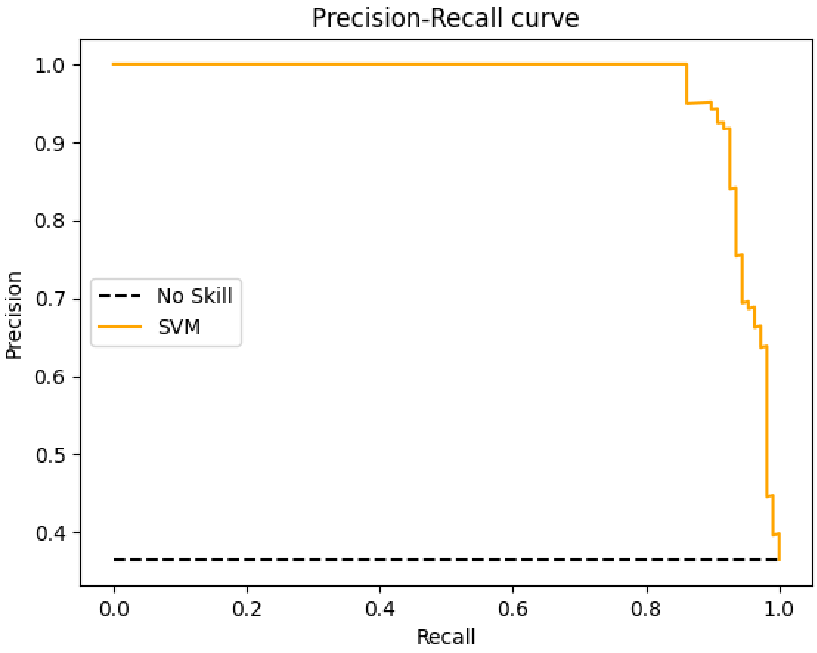

Figure 10: Precision-Recall Curve for SVM Classifier

Figure 10 shows the Precision-Recall curve for the Support Vector Machine (SVM) classifier, compared against a no-skill classifier. The SVM curve (in orange) demonstrates high precision across a wide range of recall values, remaining close to 1.0 until recall approaches its maximum. This indicates that the SVM model is highly effective at correctly identifying positive cases (fraudulent transactions), with very few false positives, especially at lower recall thresholds. As recall increases towards 1.0, there is a gradual drop in precision, which is expected as the model attempts to capture all positive cases, including some that may be incorrectly classified. The no-skill classifier, represented by the dashed black line, maintains a constant low precision, highlighting the superior performance of the SVM model. Overall, the high area under the Precision-Recall curve suggests that the SVM classifier is particularly well-suited for imbalanced datasets, such as fraud detection, where correctly identifying the minority class is crucial.

Discussion

The results of this AI-based fraud detection system demonstrate promising performance in identifying fraudulent credit card transactions, with a clear emphasis on both its strengths and areas for improvement. The classification report reveals high accuracy in detecting legitimate transactions (Class 0) but highlights challenges with detecting fraudulent transactions (Class 1), which, despite a solid precision, recall, and F1-score of 75%, remain relatively underrepresented in the dataset. The significant class imbalance between legitimate and fraudulent transactions necessitates the use of specialized techniques like oversampling or under sampling to enhance the model's sensitivity to the minority class. The decision tree model's visualization provides insights into the decision-making process, although its complexity raises concerns about overfitting. Pruning or adopting ensemble methods like Random Forest or XGBoost could help mitigate this risk and improve generalization to unseen data. Additionally, the KNN error analysis emphasizes the importance of selecting an optimal K value to balance model bias and variance, ensuring the model's robustness. The confusion matrices for both the decision tree and SVM classifiers reveal that while the model excels in identifying non-fraudulent transactions, false negatives remain a critical issue.

These missed fraud cases pose a potential financial risk, highlighting the need for further tuning to reduce false negatives and improve recall. The ROC and Precision-Recall curves for the SVM classifier show excellent performance, particularly for the minority class, reinforcing its suitability for fraud detection tasks. Overall, the system demonstrates great potential but requires additional refinements for enhanced fraud detection accuracy.

5. CONCLUSION

This study introduces a robust and comprehensive approach to developing an AI-based real-time fraud detection system tailored for credit card transactions. Utilizing a combination of machine learning algorithms—including K-Nearest Neighbors (KNN), Support Vector Machines (SVM), Decision Trees, and Logistic Regression—the system is designed to identify anomalies indicative of fraudulent behavior with high accuracy. The research leverages a publicly available, highly imbalanced Kaggle credit card dataset, which presents a realistic challenge due to the scarcity of fraud cases compared to legitimate ones. Among the models used, the Decision Tree classifier stood out by offering interpretable, rule-based insights into how decisions are made. Evaluation was conducted using essential performance metrics such as Precision, Recall, F1-Score, and ROC-AUC, ensuring a thorough analysis of model effectiveness. The system demonstrated strong overall performance, particularly in accurately identifying legitimate transactions. While detection of fraudulent transactions yielded reasonably high precision and recall, the presence of false negatives highlighted the inherent difficulty of working with imbalanced data. Visualizations including feature distribution plots, classification reports, and confusion matrices were critical in diagnosing model behavior and identifying areas for improvement. To boost performance further, particularly in real-time deployments, future work should explore advanced ensemble methods like Random Forest and XGBoost, as well as data balancing techniques such as SMOTE. Integration of cost-sensitive learning and real-time data streaming frameworks will be essential for scalability and adaptability. this study underscores the potential of AI to revolutionize financial security by enabling intelligent, real-time fraud detection systems, significantly reducing financial risk in the digital economy.

Future Scope

The future scope of AI-based real-time fraud detection systems in credit card transactions is vast and promising, with numerous avenues for further development and enhancement. As fraudsters continuously evolve their techniques, future systems must become increasingly adaptive and intelligent. One key area of future research is the integration of advanced deep learning architectures, such as Graph Neural Networks (GNNs) and Transformers, which can capture complex relationships and temporal patterns in transactional data more effectively than traditional models. Another promising direction is the incorporation of federated learning, which allows models to be trained on decentralized data across multiple financial institutions without compromising customer privacy. This approach can improve the generalizability of fraud detection models while adhering to stringent data protection regulations. Additionally, combining AI with blockchain technology can enhance transparency and traceability, providing an immutable ledger of transactions that is resistant to tampering. The use of real-time streaming data processing frameworks such as Apache Kafka and Apache Flink will also be crucial in improving the speed and responsiveness of fraud detection systems, enabling immediate actions on suspicious transactions. Moreover, techniques like adaptive anomaly detection, unsupervised learning, and self-supervised learning can help systems discover previously unseen fraud patterns without relying solely on labeled data. improving interpretability and explainability of AI models will be essential for gaining trust from users and regulatory bodies. Future work should also address challenges like reducing false positives and handling data imbalance through advanced resampling or cost-sensitive methods. These advancements will make fraud detection systems more robust, secure, and scalable.

REFERENCES

1. A. Zainal, "Role of Artificial Intelligence and Big Data Technologies in Enhancing Anomaly Detection and Fraud Prevention in Digital Banking Systems," 2023.
2. D. Beeram and J. Logeshwaran, "Real-Time Transaction Classification and Fraud Detection in Banking using AI and Advanced Data Processing," in 2024 4th International Conference on Ubiquitous Computing and Intelligent Information Systems (ICUIS), Gobichettipalayam, India: IEEE, Dec. 2024, pp. 301–306. doi: 10.1109/ICUIS64676.2024.10866509.
3. M. Baker, A. Y. Fard, H. Althuwaini, and M. B. Shadmand, "Real-Time AI-Based Anomaly Detection and Classification in Power Electronics Dominated Grids," IEEE J. Emerg. Sel. Top. Ind. Electron., vol. 4, no. 2, pp. 549–559, Apr. 2023, doi: 10.1109/JESTIE.2022.3227005.
4. D. K. -, "Real-Time AI Systems for Fraud Detection and Credit Risk Management: A Framework for Financial Institutions," IJSAT, vol. 16, no. 1, p. 2974, Mar. 2025, doi: 10.71097/IJSAT.v16.i1.2974.
5. H. K. Sriram, "Leveraging AI and Machine Learning for Enhancing Secure Payment Processing: A Study on Generative AI Applications in Real-Time Fraud Detection and Prevention," SSRN Journal, 2025, doi: 10.2139/ssrn.5203586.
6. N. A. Faisal et al., "Fraud Detection In Banking Leveraging Ai To Identify And Prevent Fraudulent Activities In Real-Time," NHJ, vol. 1, no. 01, pp. 181–197, Oct. 2024, doi: 10.70008/jmldeds.v1i01.53.
7. M. Thilagavathi, R. Saranyadevi, N. Vijayakumar, K. Selvi, L. Anitha, and K. Sudharson, "AI-Driven Fraud Detection in Financial Transactions with Graph Neural Networks and Anomaly Detection," in 2024 International Conference on Science Technology Engineering and Management (ICSTEM), Coimbatore, India: IEEE, Apr. 2024, pp. 1–6. doi: 10.1109/ICSTEM61137.2024.10560838.
8. R. A. Alzahrani and M. Aljabri, "AI-Based Techniques for Ad Click Fraud Detection and Prevention: Review and Research Directions," JSAN, vol. 12, no. 1, p. 4, Dec. 2022, doi: 10.3390/jsan12010004.
9. C. J. Zhang, A. Q. Gill, B. Liu, and M. J. Anwar, "AI-based Identity Fraud Detection: A Systematic Review," 2025, arXiv. doi: 10.48550/ARXIV.2501.09239.