JOURNAL OF MANAGEMENT AND BUSINESS SOLUTIONS

A Model for Detecting Fraud in Banking Transactions Based on Artificial Intelligence: Gambling-Related Transactions

1. Mehdi. Savadkouhi Aghamaleki : PhD Student, Department of Information Technology Management, Qa.C., Islamic Azad University, Qazvin, Iran
2. Mohammad Reza. Sanaei*: Department of Information Technology Management, Qa.C., Islamic Azad University, Qazvin, Iran
3. Soudeh. Bakhshandeh: Department of Computer Engineering, ET.C., Islamic Azad University, Tehran, Iran
3. Yashar. Bani Hashem: Department of Electrical and Computer Engineering, Buin Zahra Higher Education Technical and Engineering Center, Qazvin, Iran

*corresponding author's email: mohamadrezasanaei@gmail.com

ABSTRACT

Banking fraud is one of the major challenges that can have significant economic consequences for society. The aim of this study was to classify and detect gambling-related activities using banking transaction data. The dataset consisted of 16,764 banking transactions collected between 2023 and 2024, belonging to 1,857 distinct bank card numbers, with the deposit status column considered as the target label. Following data preprocessing procedures, including data cleaning, normalization, categorical-to-numerical conversion, and standard scaling, statistical features were extracted from two primary variables, resulting in the generation of 44 new features. Due to the high dimensionality of the feature space and the limited dataset size, three dimensionality reduction approaches based on statistical feature-selection tests were employed. The dataset was divided into training and testing subsets using an 80:20 ratio. In the baseline model, the K-Nearest Neighbors (KNN), Support Vector Machine (SVM), Random Forest, Extreme Gradient Boosting (XGBoost), Multilayer Perceptron (MLP), and Convolutional Neural Network (CNN) algorithms were evaluated using the original feature set. The proposed model was an ensemble learning framework that combined XGBoost and Random Forest classifiers through a soft-voting mechanism. Model hyperparameters were optimized using a greedy search strategy. The results demonstrated that the statistical test-based feature selection method achieved the best performance among the dimensionality reduction techniques, with an accuracy of 80.38%. Under this feature-selection approach, the proposed ensemble learning model also achieved an accuracy of 80.38%, representing an improvement of more than 6% compared with the best-performing baseline model (XGBoost with an accuracy of 73.92%). Furthermore, deep neural networks exhibited overfitting throughout all experimental stages, indicating that the available data volume was insufficient for deep learning approaches. Overall, targeted feature engineering, statistical test-based dimensionality reduction, and ensemble learning constitute an effective approach for classifying banking transactions when data availability is limited.

Keywords: Banking transactions, Dimensionality reduction, Ensemble learning, XGBoost, Random Forest, Autoencoder.

Introduction

The rapid expansion of digital banking services, electronic payment systems, and online financial transactions has transformed the global financial landscape. While these technological advancements have significantly improved accessibility, efficiency, and customer convenience, they have simultaneously increased the vulnerability

Article history:
Received 25 February 2026
Revised 10 June 2026
Accepted 15 June 2026
Initial Publish 26 June 2026
Published online 01 February 2027

How to cite this article:
Savadkouhi Aghamaleki, M., Sanaei, M. R., Bakhshandeh, S., & Bani Hashem, Y. (2027). A Model for Detecting Fraud in Banking Transactions Based on Artificial Intelligence: Gambling-Related Transactions. *Journal of Management and Business Solutions, 5*(2), 1-17. https://doi.org/10.61838/jmbs.364

© 2027 the authors. This is an open access article under the terms of the Creative Commons Attribution-NonCommercial 4.0 International (CC BY-NC 4.0) License.

of financial institutions to fraudulent activities. Financial fraud, particularly fraud associated with banking transactions and credit card operations, has become one of the most pressing challenges facing the banking industry. Fraudulent transactions generate substantial financial losses for banks, customers, and regulatory institutions, undermine trust in financial systems, and impose significant operational and reputational costs on financial organizations. Consequently, the development of intelligent and effective fraud detection systems has emerged as a strategic priority for financial institutions worldwide (1, 2).

The detection of fraudulent banking activities is inherently challenging because fraudulent transactions typically constitute only a very small proportion of the total transaction volume. This class imbalance problem makes it difficult for conventional statistical and machine learning models to identify fraudulent patterns accurately. Moreover, fraudsters continuously adapt their behaviors and employ increasingly sophisticated strategies to evade detection systems, resulting in dynamic and evolving fraud patterns. As a result, fraud detection models must be capable of identifying subtle anomalies while maintaining high levels of precision and minimizing false-positive rates. Traditional rule-based systems, which were widely used in earlier banking environments, often struggle to cope with these challenges because they rely on predefined patterns and cannot easily adapt to new fraud behaviors (2, 3).

Artificial intelligence and machine learning technologies have fundamentally changed the way financial fraud detection is approached. Unlike rule-based systems, machine learning algorithms can automatically learn complex patterns from historical transaction data and identify suspicious activities with greater flexibility and accuracy. Recent research has demonstrated that machine learning methods outperform traditional fraud detection approaches by leveraging large-scale transaction datasets and discovering hidden relationships among variables. These capabilities have encouraged researchers and practitioners to investigate a wide range of machine learning and deep learning techniques for fraud detection applications (1, 4).

Among machine learning techniques, supervised classification algorithms have received substantial attention. Algorithms such as Random Forest, Support Vector Machine, K-Nearest Neighbors, and Extreme Gradient Boosting (XGBoost) have been extensively applied to fraud detection problems due to their ability to classify transactions based on previously observed patterns. Comparative investigations have revealed that ensemble-based methods frequently outperform individual classifiers because they combine multiple decision mechanisms and reduce prediction variance. Studies comparing Random Forest and XGBoost have shown that both algorithms are highly effective for fraud classification, although performance often depends on feature quality, dataset characteristics, and parameter optimization strategies (5, 6). Furthermore, comprehensive evaluations of traditional and ensemble learning methods have highlighted the advantages of ensemble approaches in improving detection accuracy and robustness under varying fraud scenarios (4).

The emergence of deep learning has further expanded the capabilities of fraud detection systems. Deep neural networks can automatically extract hierarchical representations from large volumes of transactional data and capture nonlinear relationships that may remain hidden from conventional machine learning methods. Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs), Long Short-Term Memory (LSTM) networks, and hybrid deep learning architectures have been increasingly employed to model complex transaction behaviors. Research has shown that hybrid CNN-LSTM frameworks with attention mechanisms can improve fraud detection performance by simultaneously capturing spatial and temporal transaction patterns (7). Similarly, advanced hybrid architectures incorporating synthetic oversampling techniques, autoencoders, CNNs, and attention mechanisms have demonstrated promising results in addressing class imbalance and enhancing fraud detection accuracy (8).

Despite their predictive power, deep learning models face several challenges when applied to banking fraud detection. One major limitation involves the substantial volume of labeled data required for effective training. Financial institutions often encounter difficulties in obtaining sufficiently large and balanced fraud datasets because fraudulent events are relatively rare. Furthermore, deep neural networks are prone to overfitting when trained on limited datasets, reducing their ability to generalize to unseen transactions. To address these limitations, researchers have proposed hybrid frameworks that combine deep learning with data augmentation and feature extraction techniques. For example, generative adversarial networks have been utilized to generate synthetic fraudulent samples and improve model performance under severe class imbalance conditions (9). Additionally, semi-supervised Bayesian generative adversarial models have recently been introduced to incorporate uncertainty estimation into fraud detection processes, thereby enhancing decision reliability in complex financial environments (10).

Feature engineering represents another critical component of successful fraud detection systems. The effectiveness of machine learning models depends heavily on the quality and relevance of input features. Raw transaction records often contain limited information regarding the behavioral characteristics of users and fraudsters. Therefore, researchers have increasingly focused on extracting statistical, temporal, and behavioral features that capture transaction dynamics more effectively. Innovative feature engineering methodologies have demonstrated substantial improvements in fraud classification performance. For instance, a novel feature engineering framework based on deep learning architectures showed that carefully designed statistical and behavioral features significantly enhance fraud detection capabilities compared with using raw transactional attributes alone (11).

Behavioral analysis has also emerged as a powerful approach for anomaly detection across diverse domains. The fundamental assumption underlying behavioral analytics is that fraudulent entities exhibit behavioral patterns that differ from those of legitimate users. By modeling transaction behavior over time, anomaly detection systems can identify deviations that may indicate fraudulent activities. Research on human behavioral pattern analysis has demonstrated the effectiveness of behavior-based anomaly detection systems in identifying unusual activities and detecting suspicious events in complex environments (12). These findings suggest that behavioral feature extraction may provide valuable insights for banking fraud detection, particularly when direct indicators of fraud are unavailable.

In recent years, graph-based learning techniques have attracted significant attention within the fraud detection literature. Financial transactions naturally form interconnected networks involving customers, accounts, merchants, and institutions. Graph Neural Networks (GNNs) can exploit these relationships by modeling transaction networks and identifying suspicious structural patterns. Studies employing graph-based architectures have reported substantial improvements in fraud detection accuracy, particularly in cases involving organized fraud schemes and hidden transactional relationships (13). Similarly, the integration of graph neural networks with autoencoder architectures has been shown to enhance real-time fraud prevention by capturing both relational and transactional information simultaneously (14).

The increasing complexity of fraud patterns has encouraged the development of hybrid and ensemble models that combine the strengths of multiple analytical approaches. Hybrid frameworks integrating optimization algorithms, machine learning classifiers, temporal analysis methods, and deep learning architectures have demonstrated promising outcomes in fraud detection tasks. Research combining GRFO optimization, K-Nearest Neighbors

classification, temporal analytics, and LSTM networks has highlighted the benefits of integrating complementary methodologies to improve predictive performance (15). Likewise, studies incorporating artificial intelligence and quantum intelligence have suggested that hybrid intelligent systems may represent the next generation of fraud detection technologies capable of addressing increasingly sophisticated financial crimes (16).

Another important trend in contemporary fraud analytics is the growing emphasis on explainability and transparency. Financial institutions operate within highly regulated environments where decision-making processes must often be interpretable and auditable. While advanced machine learning and deep learning models can achieve impressive predictive accuracy, their black-box nature may limit practical adoption in regulated banking contexts. Consequently, explainable artificial intelligence (XAI) has become an essential research area. Methods such as SHAP-based interpretation frameworks provide insights into model predictions and enable financial institutions to understand the factors contributing to fraud classification decisions (17). Moreover, explainable AI approaches have been recognized as critical tools for bridging the gap between predictive fraud models and regulatory requirements in auditing and risk management applications (18).

Although substantial progress has been made in the development of fraud detection technologies, several challenges remain unresolved. Many advanced deep learning and graph-based approaches require large-scale datasets and extensive computational resources, limiting their applicability in organizations with constrained data availability. In addition, high-dimensional feature spaces can increase computational complexity and contribute to model overfitting. Therefore, dimensionality reduction techniques such as statistical feature selection, principal component analysis, and autoencoder-based representation learning have gained importance as methods for reducing feature redundancy while preserving essential information. These approaches can improve computational efficiency, enhance model generalization, and facilitate the identification of the most informative transaction characteristics (3, 6).

Recent comparative investigations have further emphasized that no single algorithm consistently outperforms all others across different fraud detection scenarios. The effectiveness of a model depends on multiple factors, including dataset characteristics, feature engineering strategies, class imbalance treatment, dimensionality reduction methods, and parameter optimization procedures. Extensive experimental comparisons have shown that ensemble learning approaches frequently achieve superior and more stable performance across diverse banking and credit fraud datasets because they leverage the complementary strengths of individual classifiers (1, 2). Consequently, the integration of carefully engineered features, effective dimensionality reduction strategies, and ensemble learning algorithms represents a promising direction for enhancing fraud detection performance in real-world banking environments.

Given the increasing prevalence of banking fraud, the limitations of existing approaches under constrained data conditions, and the demonstrated potential of feature engineering, dimensionality reduction, and ensemble learning techniques, the present study aims to develop and evaluate an artificial intelligence-based model for detecting fraudulent gambling-related banking transactions through statistical feature extraction, dimensionality reduction methods, and an ensemble learning framework combining XGBoost and Random Forest classifiers.

Methods and Materials

The data collected in this study related to banking transactions from 2023 to 2024. After cleaning, the dataset included 16,764 observations belonging to 1,857 distinct card numbers, and the objective was to classify fraudulent

individuals based on these card numbers. Of the entire dataset, 80% (1,485 card numbers) were allocated to model training and 20% (372 card numbers) were allocated to model testing.

The initial features of the dataset included transaction date (one day earlier), number of card-to-card transactions, report date, total amount of card-to-card transactions, natural-person status, transactions with identical amounts, low current account balance, account opening date, marital status, birth year, education, occupation, gender, deposit type, branch name, geographical region, and deposit status. Duplicate and equally weighted features, as well as the “report date” feature, were removed because they were not applicable. The “deposit status” column, with two values (open/closed), was considered the target label. The preprocessing steps included, in order, examining each column individually and relocating values that had been placed in incorrect columns, removing rows with missing data, standardizing values with different spellings, and converting qualitative categorical data, such as gender, education, and deposit type, into numerical values. Finally, to eliminate the effect of different feature scales, the values of each column were standardized using standard scaling, that is, subtracting the mean and dividing by the standard deviation. After these steps, the final dataset consisted of 12 columns and 16,764 rows, in which all features, except date variables, had a mean of zero and a variance of one.

To identify abnormal patterns and anomalies, new statistical features were extracted from two base features, namely “number of card-to-card transactions” and “total amount of card-to-card transactions.” These features included mean, variance, standard deviation, skewness, kurtosis, median, maximum, minimum, Gini coefficient, and transaction amount growth rate. Ultimately, 44 new features were generated.

Due to the large number of features and the limited data volume, three dimensionality reduction methods were applied separately. In the first method, namely statistical test-based feature selection, a statistical test was performed for each feature to examine the significance of its relationship with the target variable, and features with a probability value lower than 0.05 were selected, resulting in the retention of 20 final features. In the second method, namely principal component analysis, a linear transformation was applied, and a threshold of 0.95 was set for explained variance; the selected components explained 95% of the total variance in the data, and 22 principal components were selected as new features. In the third method, namely an autoencoder neural network, a neural network with an encoder–decoder architecture was designed; after training, only the encoder component was used to generate 25 compressed and nonlinear features.

In the baseline model, the K-Nearest Neighbors, Support Vector Machine, Random Forest, XGBoost, Multilayer Perceptron, and Convolutional Neural Network algorithms were evaluated using only the initial features. In the proposed model, an ensemble learning approach was used by combining two algorithms, XGBoost and Random Forest. The decision-making method in this approach was soft voting, in which each algorithm calculated the probability of each sample belonging to each class, and the average of the probabilities was considered the final output.

The algorithm parameters were determined using a greedy search method. The optimal parameters of the XGBoost algorithm included a learning rate of 0.3, a maximum tree depth of 11, 310 trees, a Lasso regularization value of 1.4, and a Ridge regularization value of 1.3. The optimal parameters of the Random Forest algorithm included 80 to 90 trees, depending on the dimensionality reduction method, and a minimum number of samples per node equal to 4 or 5. Model performance was evaluated using accuracy, precision, recall, F1-score, and area under the ROC curve. The results were reported in the form of tables, confusion matrices, and diagrams.

Findings and Results

As stated in the previous sections, the objective of this study was to classify banking transactions and identify fraudulent account numbers in banking data. The quantitative data are presented in Table 1. In addition, the distribution plots of the quantitative data and the boxplots based on deposit status are presented subsequently.

Table 1. Characteristics of the Research Dataset

Statistic	Number of Card-to-Card Transactions	Total Amount of Card-to-Card Transactions	Transactions with Identical Amounts	Low Current Account Balance	Account Opening Year	Marital Status	Birth Year	Gender	Deposit Status
Count	16,764	16,764	16,764	16,764	16,764	16,764	16,764	16,764	16,764
Mean	66.1	3.4×10^8	0.04	0.65	2020.52	0.81	1991.71	0.63	0.47
Standard deviation	117.2	3.7×10^8	0.21	0.47	4.09	0.38	11.39	0.48	0.49
Minimum	21	3.7×10^4	0	0	1998	0	1942	0	0
First quartile	31	9.5×10^7	0	0	2018	1	1985	0	0
Median	47	3.0×10^8	0	1	2023	1	1995	1	0
Third quartile	76	4.4×10^8	0	1	2023	1	2001	1	1
Maximum	8,024	1.5×10^{10}	1	1	2031	1	2020	1	1

To reduce feature dimensionality, as shown in Figure 1, effective features were extracted from the raw data. The data related to each card number were converted into an aggregated feature. Using statistical features such as mean, variance, skewness, and kurtosis, 44 new features were obtained from the two base features of transaction number and transaction amount. The importance of these features was displayed using the Random Forest algorithm on the training data. The results showed that many of the extracted features had greater importance than the initial features.

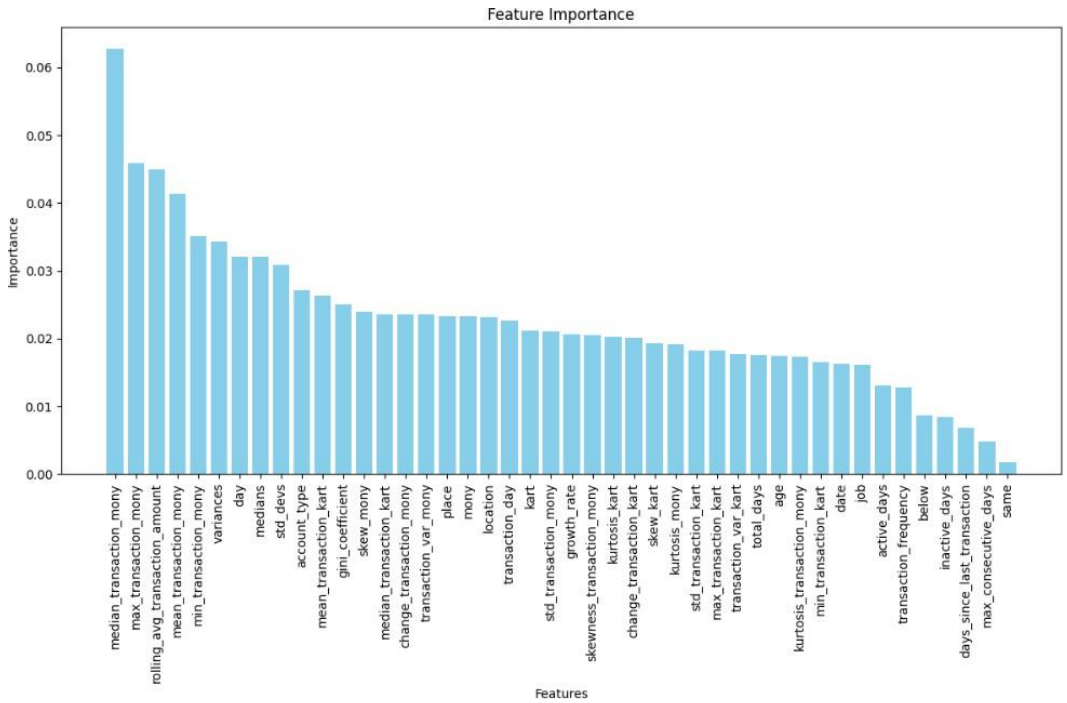

Figure 1. Importance Percentage of the Generated Features

During the statistical test-based feature selection analysis, the dimensionality of the features was reduced because of the large number of extracted features. A statistical test-based feature selection method was used, and a statistical test was performed for each feature. Features whose probability values were not lower than 0.05 were removed. In addition, the feature “low current account balance” was empirically added to the feature set. Finally, after dimensionality reduction, 20 final features were selected, and the machine learning algorithms were trained using them (Table 2). The results showed that most of the final features were extracted from the two features of the number and amount of card-to-card transactions, emphasizing the importance of these two features.

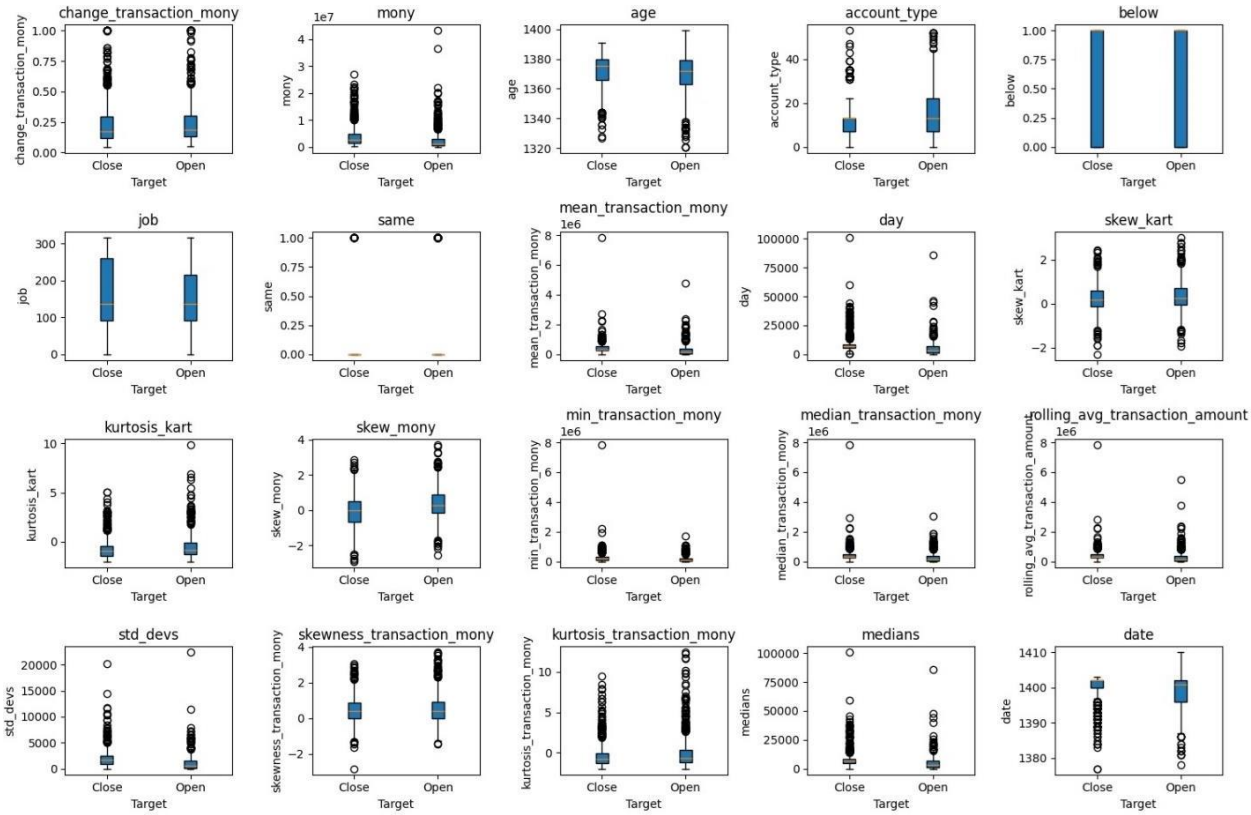

Figure 2. Boxplot of the Extracted Features

Table 2. Extracted Features and Their Probability Values

Feature	Probability Value	Feature	Probability Value
Ratio of the maximum transaction amount to the total amount of card-to-card transactions	0.001	Mean total amount of card-to-card transactions	1.3×10^{-15}
Total amount of card-to-card transactions	1.7×10^{-14}	Ratio of the total amount of card-to-card transactions to the number of card-to-card transactions	1.01×10^{-25}
Birth year	1.8×10^{-6}	Skewness value of the number of card-to-card transactions	0.0008
Deposit type	8.2×10^{-27}	Kurtosis value of the number of card-to-card transactions	3.2×10^{-8}
Low current account balance	0.06	Skewness value of the total amount of card-to-card transactions	4.6×10^{-20}
Occupation	0.004	Minimum total amount of card-to-card transactions	2.07×10^{-15}
Transactions with identical amounts	1.3×10^{-9}	Median total amount of card-to-card transactions	2.1×10^{-20}
Account opening year	9.9×10^{-14}	Mean over a specified day-window for the total amount of card-to-card transactions	3.1×10^{-14}

Median ratio of the total amount of card-to-card transactions to the number of card-to-card transactions	1.1×10^{-26}	Standard deviation of the ratio of the total amount of card-to-card transactions to the number of card-to-card transactions	1.9×10^{-31}
Kurtosis of the ratio of the total amount of card-to-card transactions to the number of card-to-card transactions	2.6×10^{-6}	Skewness of the ratio of the total amount of card-to-card transactions to the number of card-to-card transactions	0.004

Feature dimensionality reduction using principal component analysis was applied with a threshold of 0.95. The components were selected in such a way that 95% of the total variance of the data was covered. This threshold was used to create a balance between dimensionality reduction and preservation of important information. In the relevant figure, the cumulative variance plot showed that selecting the 0.95 threshold resulted in 22 final components. Thus, while important information was preserved, model complexity was reduced (Figure 3).

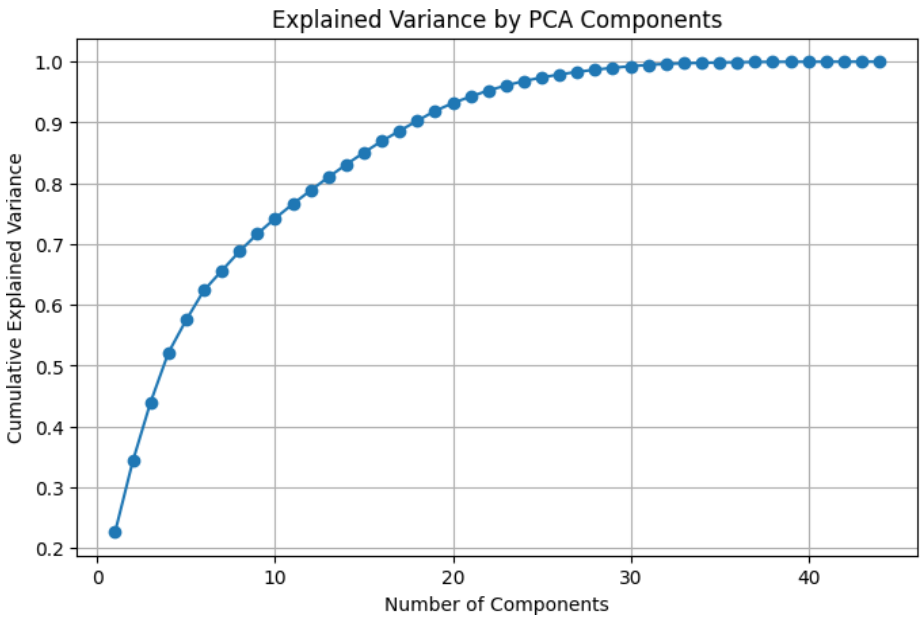

Figure 3. Cumulative Variance Plot in Principal Component Analysis

In Figure 4, the loss function plot of the trained autoencoder is displayed. Figure 5 also presents the actual values and the values reconstructed by the autoencoder, showing that the network was able to reconstruct the training and testing data appropriately. The network parameters included an input dimension of 44, one hidden layer with 32 neurons, and a bottleneck layer with 25 neurons. The Adam optimizer with a learning rate of 0.001 and the mean squared error loss function were used, and early stopping was applied to prevent overfitting. The mean squared error of the network was reported as 0.08. In addition, 20% of the data were allocated to testing and 80% to training and validation.

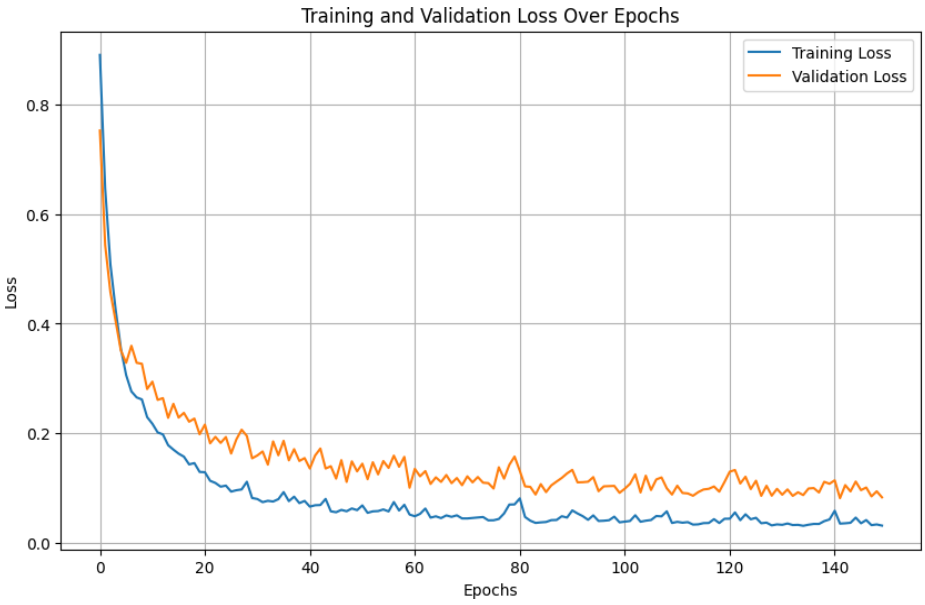

Figure 4. Loss Function Plot for the Autoencoder Network

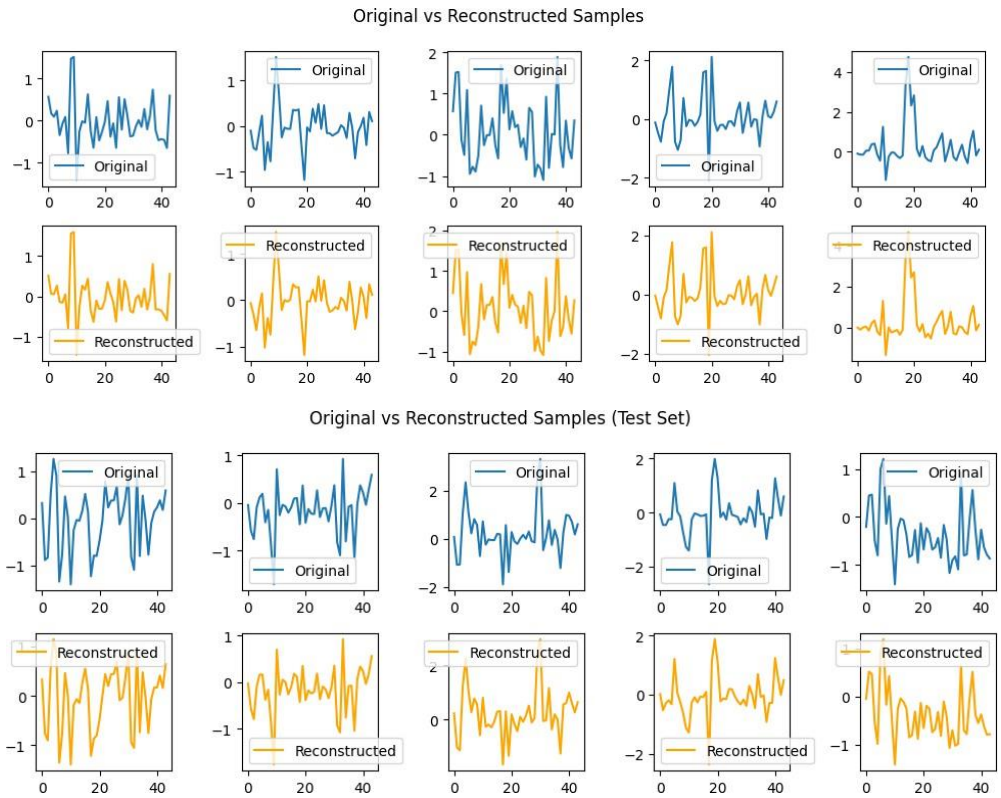

Figure 5. Actual and Reconstructed Values by the Autoencoder in the Training and Testing Datasets, Respectively

In the proposed model, an ensemble learning method combining the XGBoost and Random Forest algorithms was used. These two algorithms were also trained using the initial features. Deep learning methods, including the Multilayer Perceptron and Convolutional Neural Network, were also used for classification. The parameters of both networks included the Adam optimizer, binary cross-entropy loss function, and early stopping to prevent overfitting.

Subsequently, the accuracy table of the initial models and the confusion matrix related to each model are presented. Based on the results, the XGBoost algorithm showed the highest performance among the other algorithms. In addition, based on the error plots, it was found that the Multilayer Perceptron and Convolutional Neural Network were overfitted to the data.

Table 3. Evaluation of Different Algorithms in the Baseline Model

Algorithm	Accuracy	Precision	Recall	F1-Score	Area Under the Curve
K-Nearest Neighbors algorithm	57.8	57.85	57.61	57.37	63.88
Support Vector Machine algorithm	62.9	63.19	62.70	62.47	69.72
Random Forest algorithm	71.24	71.53	71.09	71.04	81.13
XGBoost algorithm	73.92	74.11	73.81	73.81	81.78
Multilayer Perceptron network	62.9	62.50	60.44	61.45	67.65
Convolutional Neural Network	62.63	65.03	51.1	52.73	70.64

Table 4. Confusion Matrix of Different Algorithms in the Baseline Model

Algorithm	Predicted Open Account: Actual Open Account	Predicted Open Account: Actual Closed Account	Predicted Closed Account: Actual Open Account	Predicted Closed Account: Actual Closed Account
K-Nearest Neighbors	126	93	64	89
Support Vector Machine	137	85	53	97
Ensemble learning method	150	57	40	125
Multilayer Perceptron	124	72	66	110
Convolutional Neural Network	140	89	50	93

Regarding the results obtained from the proposed model using statistical test-based dimensionality reduction, new features were first extracted. Then, to reduce model complexity, feature dimensionality was reduced using the statistical test-based method, which led to the selection of 20 features. The parameters of both algorithms were determined through greedy search. For comparison, the K-Nearest Neighbors, Support Vector Machine, Multilayer Perceptron, and Convolutional Neural Network algorithms were also employed. Finally, the accuracy table of the algorithms and the proposed model, along with the corresponding confusion matrices, was presented. According to Table 5, model performance improved across all algorithms. Moreover, the proposed model, based on the combination of XGBoost and Random Forest, achieved an accuracy of 80.38%, clearly showing the best performance across all criteria. This increase in accuracy indicates that the extracted features and the dimensionality reduction method were appropriate and were able to highlight the differences between the two groups more effectively.

Table 5. Evaluation of Different Algorithms in the Proposed Model Using Statistical Test-Based Dimensionality Reduction

Algorithm with Features and Optimized Parameters	Accuracy	Precision	Recall	F1-Score	Area Under the Curve
K-Nearest Neighbors algorithm	70.16	70.39	70.03	69.98	74.58
Support Vector Machine algorithm	72.31	72.52	72.19	72.17	75.84
Ensemble learning method (combination of XGBoost and Random Forest algorithms)	80.38	80.35	80.40	80.36	85.75
Multilayer Perceptron network	72.58	75.00	65.93	70.18	77.17
Convolutional Neural Network	70.43	70.45	68.13	69.27	76.09

Table 6. Confusion Matrix of Different Algorithms in the Proposed Model Using Statistical Test-Based Dimensionality Reduction

Algorithm	Predicted Open Account: Actual Open Account	Predicted Open Account: Actual Closed Account	Predicted Closed Account: Actual Open Account	Predicted Closed Account: Actual Closed Account
K-Nearest Neighbors	145	66	45	116
Support Vector Machine	148	61	42	121
Ensemble learning method	155	34	39	144
Multilayer Perceptron	150	62	40	120
Convolutional Neural Network	138	58	52	124

Regarding the results of the proposed model using principal component analysis for dimensionality reduction, new features were first extracted. Then, to reduce model complexity, feature dimensionality was reduced using principal component analysis, which led to the selection of 22 features. The parameters of both algorithms were determined through greedy search. Finally, the accuracy table of the algorithms and the proposed model, along with the confusion matrices, was presented (Table 7). The results showed that the performance of all algorithms improved, and the proposed model achieved an accuracy of 76.34%, which can be considered appropriate given the limited number of data. This increase in accuracy indicated that dimensionality reduction using principal component analysis was effective to some extent.

Table 7. Evaluation of Different Algorithms in the Proposed Model Using Principal Component Analysis for Dimensionality Reduction

Algorithm with Features and Optimized Parameters	Accuracy	Precision	Recall	F1-Score	Area Under the Curve
K-Nearest Neighbors algorithm	69.62	69.75	69.51	69.49	75.02
Support Vector Machine algorithm	73.92	74.22	73.79	73.77	77.04
Ensemble learning method (combination of XGBoost and Random Forest algorithms)	76.34	76.46	76.14	76.19	82.61
Multilayer Perceptron network	74.46	73.99	71.91	72.93	80.25
Convolutional Neural Network	75.27	81.16	62.92	70.89	81.73

Table 8. Confusion Matrix of Different Algorithms in the Proposed Model Using Principal Component Analysis for Dimensionality Reduction

Algorithm	Predicted Open Account: Actual Open Account	Predicted Open Account: Actual Closed Account	Predicted Closed Account: Actual Open Account	Predicted Closed Account: Actual Closed Account
K-Nearest Neighbors	142	65	48	117
Support Vector Machine	152	59	38	123
Ensemble learning method	157	51	37	127
Multilayer Perceptron	149	50	45	128
Convolutional Neural Network	168	66	26	112

Regarding the results obtained from the proposed model using autoencoder-based dimensionality reduction, new features were first extracted. Then, to reduce model complexity, feature dimensionality was reduced using the

autoencoder network, which led to the selection of 25 features. The parameters of both algorithms were determined through greedy search. For comparison, the K-Nearest Neighbors, Support Vector Machine, Multilayer Perceptron, and Convolutional Neural Network algorithms were also used. Finally, the accuracy table of the algorithms and the proposed model, along with the confusion matrices, was presented. The results showed that the performance of all algorithms improved, and the proposed model achieved an accuracy of 78.23%, indicating the suitability of dimensionality reduction using the autoencoder network and the clearer differentiation between the two groups.

Table 9. Evaluation of Different Algorithms in the Proposed Model Using Autoencoder-Based Dimensionality Reduction

Algorithm with Features and Optimized Parameters	Accuracy	Precision	Recall	F1-Score	Area Under the Curve
K-Nearest Neighbors algorithm	69.35	69.46	69.25	69.23	75.33
Support Vector Machine algorithm	71.24	71.59	71.08	71.01	75.86
Ensemble learning method (combination of XGBoost and Random Forest algorithms)	78.23	78.27	78.08	78.13	85.02
Multilayer Perceptron network	75.54	77.36	69.10	73.00	83.30
Convolutional Neural Network	72.58	73.91	66.85	70.21	79.12

Table 10. Confusion Matrix of Different Algorithms in the Proposed Model Using Autoencoder-Based Dimensionality Reduction

Algorithm	Predicted Open Account: Actual Open Account	Predicted Open Account: Actual Closed Account	Predicted Closed Account: Actual Open Account	Predicted Closed Account: Actual Closed Account
K-Nearest Neighbors	149	66	41	116
Support Vector Machine	158	45	36	133
Ensemble learning method	158	55	36	123
Multilayer Perceptron	152	59	42	119
Convolutional Neural Network	149	66	41	116

Discussion and Conclusion

The present study aimed to develop an artificial intelligence-based framework for identifying fraudulent gambling-related banking transactions through statistical feature extraction, dimensionality reduction techniques, and ensemble learning. The findings demonstrated that the proposed framework successfully improved fraud classification performance compared with conventional machine learning and deep learning approaches. More specifically, the ensemble learning model combining XGBoost and Random Forest achieved the highest classification accuracy when statistical test-based feature selection was employed, reaching an accuracy of 80.38%, which represented a substantial improvement over the best baseline model. The results provide important evidence regarding the value of targeted feature engineering and dimensionality reduction in banking fraud detection environments characterized by limited data availability.

One of the most significant findings of the study was the superior performance of the statistical test-based feature selection approach compared with principal component analysis and autoencoder-based dimensionality reduction. Statistical feature selection reduced the original feature space to 20 highly relevant variables while preserving the information most strongly associated with fraudulent behavior. This outcome suggests that identifying statistically significant features can be more effective than purely mathematical dimensionality reduction approaches when the

objective is to distinguish between fraudulent and legitimate banking transactions. The finding aligns with previous studies emphasizing the critical role of feature engineering in fraud detection. Zhang et al. demonstrated that carefully designed statistical and behavioral features substantially improve classification performance compared with raw transactional data alone (11). Similarly, Li reported that effective feature selection strategies contribute significantly to the success of ensemble fraud detection systems, particularly under conditions of class imbalance and limited data availability (3).

The results further revealed that most of the selected features originated from the number and amount of card-to-card transactions, indicating that transaction frequency and transaction volume constitute key indicators of suspicious behavior. Variables derived from skewness, kurtosis, transaction ratios, moving averages, and statistical distributions exhibited stronger predictive power than many demographic and account-related attributes. This finding is theoretically consistent with behavioral anomaly detection research, which argues that abnormal patterns often emerge through deviations in transactional behavior rather than through static customer characteristics. Choi et al. demonstrated that behavioral patterns provide valuable information for identifying anomalous activities, and the current results extend this principle to banking transaction analysis (12). The findings suggest that fraud detection systems should focus on dynamic behavioral indicators rather than relying solely on demographic information.

Another important outcome was the strong performance of the proposed ensemble learning model. The combination of XGBoost and Random Forest consistently outperformed individual machine learning algorithms across all dimensionality reduction scenarios. The ensemble model achieved 80.38% accuracy with statistical feature selection, 78.23% accuracy with autoencoder-based dimensionality reduction, and 76.34% accuracy with principal component analysis. These results support the growing consensus that ensemble learning approaches provide greater robustness and predictive capability than single classifiers. Previous research has consistently reported the effectiveness of ensemble methodologies in fraud detection contexts. Khalid et al. demonstrated that ensemble machine learning approaches outperform individual algorithms by integrating multiple decision-making processes and reducing model variance (6). Likewise, Bothra reported that ensemble models generally achieve higher stability and classification performance across diverse online fraud detection datasets (4). The findings of the present study provide additional empirical support for the superiority of ensemble learning in financial fraud detection applications.

The performance of XGBoost as the strongest baseline algorithm is also noteworthy. Among the baseline models trained on the original features, XGBoost achieved an accuracy of 73.92%, outperforming Random Forest, Support Vector Machine, K-Nearest Neighbors, and the deep learning models. This finding is consistent with previous studies comparing fraud detection algorithms. Abdullah et al. found that XGBoost frequently exceeds Random Forest performance due to its ability to model complex nonlinear relationships and iteratively reduce classification errors through gradient boosting mechanisms (5). Similarly, Mohammed reported that gradient boosting methods consistently rank among the best-performing algorithms in credit card fraud detection tasks because of their strong generalization capabilities and resistance to overfitting under moderate data conditions (2). The current findings reinforce the suitability of XGBoost as a core component of advanced fraud detection systems.

An additional contribution of the study concerns the evaluation of principal component analysis and autoencoder-based dimensionality reduction. Both methods improved classification performance compared with the baseline models, although neither matched the effectiveness of statistical feature selection. Principal component analysis

reduced the feature space to 22 components while retaining 95% of the variance, resulting in a meaningful increase in classification accuracy. This finding suggests that eliminating redundant information and reducing dimensionality can improve model efficiency and predictive performance. Previous studies have highlighted the importance of dimensionality reduction in fraud detection systems because high-dimensional data often introduce noise and increase computational complexity (1, 3). Therefore, the observed improvements following principal component analysis support the broader literature regarding the benefits of feature-space optimization.

The autoencoder-based dimensionality reduction approach also produced encouraging results. The proposed ensemble model achieved an accuracy of 78.23% when trained on the compressed feature representation generated by the autoencoder. This outcome suggests that autoencoders are capable of capturing meaningful nonlinear structures within banking transaction data. The findings align with recent research demonstrating the utility of autoencoder architectures for financial fraud detection. Alarfaj and Shahzadi showed that autoencoders can effectively identify latent patterns associated with fraudulent behavior while reducing feature complexity (14). Similarly, Hasan et al. reported that autoencoder-based feature representations improve deep learning performance by generating more informative and compact data representations (8). The current results support these conclusions while also indicating that statistical feature selection may remain preferable when data availability is limited.

The behavior of the deep learning models provides another important insight. Although deep learning methods have achieved remarkable success in many fraud detection applications, the Multilayer Perceptron and Convolutional Neural Network models in this study exhibited clear signs of overfitting. Training performance increased while generalization performance remained relatively limited, suggesting that the available dataset size was insufficient to fully exploit the representational power of deep neural networks. This observation is consistent with existing literature indicating that deep learning models require large-scale datasets to achieve optimal performance. Ahmed et al. reported substantial performance gains for hybrid CNN-LSTM architectures when trained on extensive fraud datasets containing rich temporal information (7). Likewise, Hasan et al. demonstrated the effectiveness of deep architectures when combined with synthetic oversampling and attention mechanisms in large-scale environments (8). The present findings indicate that such approaches may be less effective when data volume is constrained.

The observed limitations of deep learning models also highlight the importance of addressing class imbalance and data scarcity in fraud detection research. Recent studies have proposed generative adversarial networks and synthetic data generation techniques as potential solutions. Mienye and Swart demonstrated that generative adversarial networks can improve fraud detection performance by generating synthetic fraudulent transactions and increasing class balance (9). Furthermore, Hirnschall introduced semi-supervised Bayesian GAN architectures capable of providing uncertainty-aware fraud detection under limited supervision conditions (10). While such approaches were beyond the scope of the current investigation, the findings suggest that integrating synthetic data generation techniques may further improve performance in future implementations.

The results additionally support recent developments involving hybrid and advanced intelligent systems. Research integrating temporal analytics, optimization methods, and deep learning architectures has demonstrated promising improvements in fraud detection performance (15). Likewise, hybrid artificial and quantum intelligence systems have been proposed as future directions for highly adaptive fraud detection environments (16). Although the present study employed a comparatively simpler ensemble structure, the significant performance gains

achieved through the combination of XGBoost and Random Forest suggest that hybridization remains a highly effective strategy for enhancing fraud detection capabilities.

Another noteworthy implication concerns model interpretability. Financial institutions increasingly require transparent and explainable decision-making processes to satisfy regulatory and auditing requirements. The superior performance of statistically selected features facilitates interpretability because each selected variable retains a direct relationship with observable transactional behaviors. This finding aligns with the growing emphasis on explainable artificial intelligence in financial applications. Sharma and Gupta demonstrated that SHAP-based explainability methods enhance trust and transparency in fraud detection systems by revealing the contribution of individual variables to classification decisions (17). Similarly, recent work in auditing and regulatory compliance has highlighted the importance of explainable models for bridging the gap between predictive performance and governance requirements (18). Therefore, the proposed framework offers practical advantages not only in predictive accuracy but also in interpretability.

The findings of this study also contribute to the broader literature on banking fraud detection by providing evidence that sophisticated graph-based architectures are not always necessary to achieve competitive performance. Although graph neural network approaches have demonstrated excellent results in complex transaction networks (13, 14), the current study shows that substantial improvements can be achieved through carefully engineered statistical features combined with ensemble learning techniques. This is particularly relevant for organizations that may lack the computational infrastructure or relational data required for graph-based implementations.

Overall, the findings demonstrate that targeted statistical feature extraction, appropriate dimensionality reduction, and ensemble learning can substantially improve the detection of fraudulent banking transactions associated with gambling activities. The superior performance of the proposed model confirms the importance of combining informative behavioral indicators with robust classification algorithms, particularly in environments characterized by limited data availability. The study therefore provides both theoretical and practical evidence supporting the integration of feature engineering and ensemble machine learning for effective banking fraud detection.

This study has several limitations that should be considered when interpreting the findings. First, the dataset was obtained from a single banking environment, which may limit the generalizability of the results to other financial institutions, geographical regions, or transaction ecosystems. Second, the number of fraudulent accounts was relatively limited, creating an imbalanced classification problem that may have influenced model learning. Third, the study focused primarily on structured transactional features and did not incorporate network-based relationships among accounts, merchants, or customers. Fourth, the temporal evolution of fraudulent behavior was not explicitly modeled, which may restrict the model's ability to detect emerging fraud patterns. Finally, although several machine learning and deep learning algorithms were evaluated, more advanced architectures and synthetic data generation techniques were not included in the comparison.

Future studies should evaluate the proposed framework using larger and more diverse datasets obtained from multiple financial institutions to enhance external validity. Researchers may also investigate the integration of graph neural networks, temporal sequence models, and transformer-based architectures to capture more complex transactional relationships. The incorporation of synthetic data generation methods, such as generative adversarial networks, could help address class imbalance and improve model generalization. Future research should further explore explainable artificial intelligence techniques to improve transparency and regulatory compliance.

Comparative studies examining real-time fraud detection systems and streaming transaction data environments would also provide valuable insights into the operational deployment of fraud detection models.

Financial institutions can improve fraud detection effectiveness by emphasizing behavioral transaction indicators rather than relying solely on demographic information. Banks should invest in feature engineering processes that extract statistical and temporal patterns from transaction histories. Ensemble learning models combining complementary algorithms can be deployed to improve predictive performance while maintaining computational efficiency. Organizations should regularly update fraud detection systems to adapt to evolving fraud strategies and incorporate continuous monitoring mechanisms. Finally, integrating interpretable machine learning frameworks into operational fraud detection systems can enhance decision transparency, facilitate regulatory compliance, and strengthen stakeholder confidence in automated fraud prevention processes.

Acknowledgments

We would like to express our appreciation and gratitude to all those who helped us carrying out this study.

Authors' Contributions

All authors equally contributed to this study.

Declaration of Interest

The authors of this article declared no conflict of interest.

Ethical Considerations

All ethical principles were adheried in conducting and writing this article.

Transparency of Data

In accordance with the principles of transparency and open research, we declare that all data and materials used in this study are available upon request.

Funding

This research was carried out independently with personal funding and without the financial support of any governmental or private institution or organization.

References

1. Yapp EKY, Yeh HY. An Extensive Experimental Comparison of Machine and Deep Learning Methods for Credit and Bank Fraud Detection. Finance Research Letters. 2026;88:109190. doi: 10.1016/j.frl.2025.109190.
2. Mohammed M. Comparison of Various Classification Algorithms for Credit Card Fraud Detection. Taylor & Francis Book Chapter: Taylor & Francis; 2026.
3. Li W, editor Credit Card Fraud Detection: A System Based on Imbalanced Learning and Ensemble Models. Proceedings of the 2025 7th International Conference on Economic Management and Model Engineering (ICEMME 2025); 2026.
4. Bothra Y. A Comprehensive Performance Comparison of Traditional and Ensemble Machine Learning Models for Online Fraud Detection. arXiv preprint. 2025:arXiv:2509.17176.

5. Abdullah A, Khairah DU, Pangestika MW. Comparison of Random Forest and XGBoost Algorithms in Credit Card Fraud Classification. Journal of Computer Science and Technology (COSCITECH). 2025;6(3). doi: 10.37859/coscitech.v6i3.10470.

6. Khalid AR, Owoh N, Uthmani O, Ashawa M, Osamor J, Adejoh J. Enhancing Credit Card Fraud Detection: An Ensemble Machine Learning Approach. Big Data and Cognitive Computing. 2024;8(1):6. doi: 10.3390/bdcc8010006.

7. Ahmed S, Khan MA, Rahman M. Hybrid CNN-LSTM with Attention Mechanism for Robust Credit Card Fraud Detection. IEEE Access. 2025;13:114056-68. doi: 10.1109/ACCESS.2025.3583253.

8. Hasan MA, Hossain MS, Islam MR. A Hybrid Deep Learning Framework Using Synthetic Oversampling, Autoencoder, Convolutional Neural Networks, and an Attention Mechanism for Credit Card Fraud Detection. Journal of Big Data. 2026;13:Article 21. doi: 10.1186/s40537-025-01331-2.

9. Mienye ID, Swart TG. A Hybrid Deep Learning Approach with Generative Adversarial Network for Credit Card Fraud Detection. Technologies. 2024;12(10):186. doi: 10.3390/technologies12100186.

10. Hirnschall D. Semi-Supervised Bayesian GANs with Log-Signatures for Uncertainty-Aware Credit Card Fraud Detection. arXiv preprint. 2026:arXiv:2509.00931. doi: 10.3390/math13193229.

11. Zhang X, Han Y, Xu W, Wang Q. HOBA: A Novel Feature Engineering Methodology for Credit Card Fraud Detection with a Deep Learning Architecture. Information Sciences. 2023;632:1-15. doi: 10.1016/j.ins.2023.02.090.

12. Choi S, Kim C, Kang YS, Youm S. Human Behavioral Pattern Analysis-Based Anomaly Detection System in Residential Space. The Journal of Supercomputing. 2021;77:9248-65. doi: 10.1007/s11227-021-03641-7.

13. Cherif A, Ammar H, Kalkatawi M, Alshehri S, Imine A. Encoder-Decoder Graph Neural Network for Credit Card Fraud Detection. Journal of King Saud University - Computer and Information Sciences. 2024;36(3):102003. doi: 10.1016/j.jksuci.2024.102003.

14. Alarfaj FK, Shahzadi S. Enhancing Fraud Detection in Banking with Deep Learning: Graph Neural Networks and Autoencoders for Real-Time Credit Card Fraud Prevention. IEEE Access. 2024;12:145678-95.

15. Choudhary I, Kapoor H, Kaushik E. Cutting Edge Hybrid Models for Credit Card Fraud Detection: Integrating GRFO-KNN, Temporal Analysis, and LSTM Networks. In: Patnaik S, Abe JM, Nakamatsu K, Vigliarolo F, editors. Social, Ethical and Legal Aspects of Generative AI. Studies in Computational Intelligence. 1185: Springer; 2025.

16. El-Fargani A, El-Hassan M. Data-Driven Financial Fraud Detection Using Hybrid Artificial and Quantum Intelligence. Journal of Open Innovation: Technology, Market, and Complexity. 2025;11(4):100581.

17. Sharma P, Gupta R, editors. Enhancing Credit Card Fraud Detection with Explainable AI: A SHAP-Based Approach. Proceedings of 2025 International Conference on Frontiers of Information Technology (FIT 2025); 2025: IEEE.

18. Journal of R, Financial M. Explainable AI (XAI) in Auditing: Bridging the Gap between Predictive Fraud Models and Regulatory Standards. Journal of Risk and Financial Management. 2026;19(5):311. doi: 10.3390/jrfm19050311.