

An Intelligent Ensemble-Based Fraud Detection System for Banking Transactions using Machine Learning and Deep Learning

B. Annapurna

R. Aparna, Sahil Awasthi, Y. Sanjana

CSE(AI&ML) CMR Institution of Technology, Hyderabad, Telangana, India

Abstract - The rapid expansion of digital banking and online financial transactions has significantly increased fraudulent activities, resulting in substantial financial losses. Traditional rule-based fraud detection systems fail to handle highly imbalanced and large-scale banking datasets effectively. This paper proposes an intelligent ensemble-based fraud detection system using optimized machine learning and deep learning techniques. The proposed system integrates Logistic Regression, Decision Tree, Support Vector Machine, Random Forest, AdaBoost, Multilayer Perceptron, and Deep Neural Networks to classify banking transactions as legitimate or fraudulent. Data preprocessing techniques including normalization, feature extraction, and dataset balancing are applied to improve model performance. The models are evaluated using accuracy, precision, recall, and F1-score. Experimental results demonstrate that ensemble and deep learning models outperform traditional classifiers, achieving higher detection accuracy while reducing false positives. The proposed approach provides an efficient and scalable solution suitable for real-time banking fraud detection applications.

Keywords : Fraud Detection, Banking Transactions, Machine Learning, Deep Learning, Ensemble Learning, Random Forest, Support Vector Machine, Imbalanced Dataset

I. INTRODUCTION

The high pace of the digital technologies development has radically changed the banking ecosystem in the world. During the last ten years financial institutions began to move away by abandoning the old services of working in the traditional branch instead of adopting the new technologies of online banking, mobile applications, digital wallets, and real-time payment systems. This online revolution resulted in a high level of convenience to the customers, cut in operations costs and speedy financial transactions between geographical borders. Nonetheless, as much as these advantages are evident, tandem with these gains, the growth of online financial services has also brought about new avenues of frauds. Like ever, cybercriminals are taking advantage of the loopholes in the online payment systems, credit card activities, internet bank sites, and digital transfer systems and causing huge financial losses to banks and users across the globe.

Banking frauds have emerged as among the most serious security threats to the financial institutions. Examples of frauds involve unauthorized transactions, identity theft, credit card abuse, account violation, transfers made through phishing and synthetic identity fraud. The challenge presented by fraud detection is due to the responsive behavior of the frauds. The

fraudsters are always changing their tactics to avoid the prevailing security systems, and the rule-based methods of detection will not work effectively in the long run. Furthermore, today banking systems execute millions of transactions a day, and they produce big loads of information that are not visible to the human eye. The magnitude and complexity of this requires smart, automated and scalable fraud detection systems that are able to process the transaction data in a real time manner.

The primary methods of traditional fraud detection system are based on pre-set rules and threshold-based system. Such systems work on hand-designed specifications like flagging transactions that are over a certain amount or preventing frequent installment of unproductive attempted logins. Although rule based systems are easy to interpret and implement, they have serious limitations. They also need constant human oversight, are not capable of identifying previously unknown patterns of fraud, and also have a high rate of false positives. False alarm is also detrimental in that it blocks valid transactions even when it is excessive and it raises the operational costs. Since financial systems become more complex, it does not suffice to provide effective fraud prevention using the rule-based approaches.

Machine learning has turned out to be an effective tool of detecting fraud in banking systems in the recent years. The pattern of the past transactions can be learned by the algorithms of machine learning to detect minor dependencies among features that could point out to some form of fraud. In contrast to a rule-based system, machine learning models are able to adjust to novelties of new strategies of fraud as it is retrained on new datasets. Detecting fraud with the help of machine learning is often declared as a binary classification problem, in which every transaction is assigned a legitimate or a fraud label following the boundaries of decisions learned. Some of the algorithms that have been extensively used in the financial fraud detection exercises are Logistic Regression, Decision Trees, the Support Vector Machines, Random forest and boosting algorithms.

Nevertheless, fraud detection has special problems, which make the process of learning more complex. A highly unbalanced nature of banking datasets is one of the biggest challenges, as well. In actual life scenarios, fraudulent transactions constitute a very small percentage of transacted transactions with most being legitimate. Such imbalance may bias the learning algorithms to predict the dominant group, giving it a high accuracy with a low ability to detect fraud. Hence, measures of evaluation like precision, recall, and F1-score are increasingly significant instead of cumulative determination when evaluating performance of the model. The complicated and non-linear relationship among the feature of the transaction is another challenge, which may not be well defined under the simple linear model.

In order to overcome such limitations, ensemble learning and deep learning methods have been acquiring increasingly importance in fraud detection literature. Ensemble learning takes the base classifiers and produces a stronger and more stable prediction. Approaches like random forest and AdaBoost use many weak learners, which can eliminate variance and bias which in many cases lead to better generalisation. Deep learning models, however, can learn high-level representations of features created out of raw input data. Neural networks have the capability to represent complicated and non-linear patterns and interactions among transaction attributes and, therefore, it is possible to detect suspicious behavior much more accurately. Ensemble and deep learning methods can be integrated to give a promising future to creating intelligent fraud detection systems.

The study hypothesizes a smart, ensemble-driven fraud detection framework of banking transactions that incorporates combination of multiple machine learning algorithms and deep neural networks in a single pipeline of experimental platform. The system also features structured preprocessing of data and normalisation of features as well as fixed splitting of train and test sets as a way of ensuring a reliable evaluation of the model.

Several of the classifiers are applied and contrasted under the same conditions to examine their efficiency with unequal transaction data. Using ensemble mechanisms and deep learning designs, the proposed framework is expected to increase the detection accuracy and reduce false positives at a lower computational cost.

The main aim of this paper is to develop scalable and reliable model of detecting fraud that can perform tasks with large banking transactions volumes. The suggested method compares legacy classifiers to ensemble and deep learning classifiers to define the most efficient method to identify a fraudulent behavior. The paper has illustrated the importance of intelligent model integration as a way of enhancing the security system of banking through thorough performance analysis of the banking system with the right evaluation standards. In the end, this study can lead to the creation of more complex and automated fraud-detecting systems that will help to ensure privacy of online financial payments and enhance the confidence in the modern banking facilities

Besides enhancing detection performance, the given framework does not focus only on increasing detection accuracy but also focuses on flexibility and reproducibility in contemporary banking settings. With the development of technologies in the sphere of finances the fraud patterns also become more complicated and dynamic, which means that they should have the opportunity to learn and develop constantly. The fact that several learning models are incorporated in a single framework guarantees in the flexibility of experimentation and upgradeability in the future. The proposed intelligent ensemble-based system

can fill the present analytical rigor and computational efficiency with multiple benefits, as the current fraud detection problems are resolved, and the future of its practical use and research is opened. This will eventually facilitate the overall intention of establishing .

II . RELATED WORK

The identification of fraud in the banking and financial systems is not a new field of research that has been actively developed over the last few decades, especially with the advent of online banking services and services based on the digital payment platform. Methods of early detection of fraud were mostly based on a statistical method and a rule-based method. These systems would determine suspicious transactions using predefined thresholds and heuristic developed by experts. Though these methods were easy to understand and interpret, they were not flexible but could not pick up the new or changing trends of fraud. As financial data were growing in size and complexity, scholars started to consider approaches that would

automatically discover patterns based on available historical transaction data.

Table1: Comparison of Existing Fraud Detection Methods

Study Approach	Method Used	Strengths	Limitations
Data Mining Approaches	Decision Tree, Bayesian Networks	Good for pattern recognition	Struggle with highly imbalanced datasets
Machine Learning Models	Logistic Regression, SVM, KNN	High classification accuracy	High false positive rate
Ensemble Learning	Random Forest, XGBoost, AdaBoost	Improved accuracy and robustness	Higher computational cost
Deep Learning Methods	Neural Networks, CNN	Detect complex fraud patterns	Less interpretable and computationally expensive

Among the research methods that are supervised machine learning models are one of the fundamentals in fraud detection research. Logistic Regression has been extensively employed as a baseline classifier thanks to the level of simplicity and ease of comprehension. It estimates the likeliness of a transaction as fraudulent employing a straight decision boundary. Although useful in a case of linearly separable data, the Logistic Regression does not function much well with complex non-linear relationship entities that modern transaction data possesses. On the same note, Naive Bayes classifiers have been used to detect the fraud due to their probability based modeling and their computational high efficiency. But their performance is constrained by the assumption of feature independence which does not ideally cope with correlated attributes of transaction.

The use of Decision Trees and their derivatives is an enormous move towards addressing the study of frauds. These models divide the dataset in accordance to feature significance and generate naturally hierarchical decision regulations. Studies have demonstrated that Decision Trees can represent non-linear relationships and they can be interpreted relatively readily thereby being appropriate when dealing with financial problems and transparency is a required factor. Nevertheless, their independent Decision Trees are susceptible to overfitting especially when they are trained on extremely skewed data. In order to counter this drawback, ensemble learning methods like the Random Forest and Boosting algorithms were suggested. Random Forest, which was put forward by Breiman, is a subset

of decision trees created on random subsets of data and features, and so - variance is reduced and there is generalization. The boosting algorithms like AdaBoost are applied in a cascading fashion; it trains weak learners that are being constantly enhanced on the bases of selectively misclassified samples thereby enhancing the overall predictive accuracy. Single classifiers relative to these ensemble methods have always proved to be worse in fraud detection tasks.

One of the areas that fraud detection studies have given serious emphasis is managing skewed datasets. In the actual banking data, fraudulent purchase may constitute less than one percent of the total purchase. Such imbalance may lead to the tendency of models to give preference of the majority class, which makes it falsely claim high accuracy and low probability of hyperfraud. Researchers have suggested some solutions to this problem including Synthetic Minority Over-sampling Technique (SMOTE), under-sampling of the majority, under-sample and using hybrid sampling methods. Cost sensitive learning has also been studied where the costs of misclassification are weighted to favor false negative more compared to false positive. These are strategies to reduce fraud recall without making any major rise in false alarms.

In financial fraud detection, gradient boosting algorithms like XGBoost, LightGBM, and CatBoost have become increasingly popular with the increasing computational power, and with large scale datasets. Such models enhance predictive performance by using optimized gradient boosting algorithms, regularization and using an efficient tree building algorithms

XGBoost proposed scalable and regularized boosting algorithms which raised the performance on structured financial data. LightGBM also enhanced performance through splitting with histograms and leaf-wise growing to provide more training power with large data volumes. CatBoost added specific categorical feature treatment which simplified the preprocessing and made it more robust. Such enhancing models have shown high prediction accuracy problems in credit card fraud and transaction tracking.

In recent years, deep learning methods were also studied in detail. Neural networks offer a learning hierarchy feature representation and complex non-linear interactions among the attributes of transactions. MLP has also been used in undertaking fraud detection activities and they have performed better than the conventional linear classifiers. Further architectures like the Convolutional Neural Networks (CNNs) have been applied to structured transaction data through reshaping the feature vectors into a higher-dimensional form. Also, Recurrent Neural Networks (RNNs) and Long Short-Term Memory (LSTM) networks have been suggested to learn temporal dependence in

a sequence of transactional behaviour. Sequential models provide an encouraging point of view of the improved detection of fraudulent activity because the activities are usually repeated over an extended period.

The other new field of investigation in fraud detection is the graph-based learning methods. Hacker activities that constitute fraudulent transactions tend to be perpetuated in networks of accounts, devices, or entities. Graph Neural Networks (GNNs) have also been researched to realize organized fraud rings and to model the relational dependencies. Graph-based models can identify anomalous clusters so that the analysis of links between accounts goes beyond the individual transaction to include network-based fraud detection. Though excellent, these approaches demand intricate infrastructure and extensive relational data in order to achieve success.

Although machine learning and deep learning solutions have been made, there are a number of challenges in the research of fraud detection today. Most of the research has concentrated on the singular performance of algorithms, as opposed to the combination of two or more models under the same system. It is not regularly feasible to conduct comparative analysis of various algorithms under similar preprocessing conditions. Also, although advanced boosting and the deep learning model can produce high accuracy, they can be not interpretable; that is needed in the controlled financial world.

The current contribution is based on current literature through the integration of established machine learning frameworks, ensemble learning methods, and deep neural networks into a system of structured experiments. With general preprocessing and normalization and assessment measures, the paper will offer a comparative systematic comparison of several methods on imbalanced banking data. The proposed integrated approach to these two models is to fill the gap between single-model research and practical needs of fraud detection in the real-world environment and provide a balanced approach, focusing on both accuracy, scalability, and reliability of the digital banking scheme.

III. METHODOLOGY

This study developed a machine learning-based framework to detect fraudulent banking transactions using historical transaction data. The methodology consisted of data preprocessing, feature preparation, model training, and performance evaluation. Multiple classification algorithms were implemented and compared to identify the most effective model for fraud detection.

A. Dataset Description

The dataset used in this study consisted of banking transaction records containing both legitimate and fraudulent transactions. Each record included several features representing transaction characteristics such as transaction amount, time, and anonymized

attributes related to the transaction process. The dataset contained a large number of normal transactions and a relatively small number of fraudulent transactions, resulting in a highly imbalanced classification problem.

To ensure reliable model performance, the dataset was divided into training and testing sets, where 80% of the data was used for model training and 20% was reserved for testing and evaluation.

B. Data Preprocessing

Before training the machine learning models, several preprocessing steps were performed to prepare the dataset for analysis.

First, missing or inconsistent values were handled to maintain data quality. Feature scaling and normalization were applied to ensure that all numerical attributes were on a comparable scale, which improves the performance of many machine learning algorithms.

Since fraud detection datasets are typically imbalanced, where fraudulent transactions represent only a small percentage of the data, appropriate preprocessing techniques were applied to minimize the impact of class imbalance during model training.




Fig 1

C. Feature Preparation

Feature preparation was performed to select relevant attributes from the transaction dataset that contribute to fraud detection. Each feature represents a measurable property of the transaction that may help distinguish fraudulent behavior from legitimate activity.

The prepared feature set was then converted into a numerical format suitable for machine learning algorithms. The target variable was defined as a binary label indicating whether a transaction was fraudulent or legitimate.

D. Machine Learning Models

Several machine learning classification algorithms were implemented to analyze and detect fraudulent transactions. The models used in this study included:

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

- Naïve Bayes
- Multilayer Perceptron Neural Network

Each model was trained using the prepared dataset to learn patterns associated with fraudulent activities. The algorithms analyzed transaction features and generated predictions indicating whether a given transaction was likely to be fraudulent.

Table 2: Performance Comparison of Machine Learning Models

E. Model Training and Testing

The machine learning models were trained using the training dataset and evaluated using the testing dataset. During training, the models learned relationships between input features and the fraud classification label.

The training process involved adjusting model parameters to minimize classification error. After training was completed, each model was applied to the testing dataset to evaluate its ability to correctly classify unseen transactions.

F. Performance Evaluation

The performance of the machine learning models was evaluated using standard classification metrics. These metrics measured how accurately the models detected fraudulent transactions while minimizing incorrect predictions.

The following evaluation metrics were used:

- Accuracy – the overall percentage of correctly classified transactions
- Precision – the proportion of predicted fraud cases that were actually fraudulent
- Recall – the proportion of actual fraud cases correctly detected
- F1-score – the harmonic mean of precision and recall

These metrics provide a comprehensive evaluation of model performance, particularly for imbalanced datasets where accuracy alone may not provide reliable insights.

IV . RESULT ANALYSIS

Several machine and deep learning models were also trained and functionality of the proposed system to detect fraud was evaluated with help of the testing dataset. The system was divided into sensible assessment of 80 percent training information and 20 percent checking information. The special attention was also paid to precision due to the measurement of recall and F1-score, as well as the accuracy bearing, considering the fact that the dataset was highly imbalanced.

The systematics of each of the implemented algorithms of Logistic Regression, Naive Bayes, Decision Tree, Support Vector Machine (SVM), Random Forest, AdaBoost, Multilayer Perceptron (MLP) and Deep Neural Network were compared.

The metrics that became established to be used in the evaluation are as below:

Measures: General prediction accuracy.

Precision: Refers to instances of predicted frauds which actually translated to a fraud.

Recall: Determines the actual occurrences of the cases of frauds that have been correctly identified.

Algorithm	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
Logistic Regression	83.6	84.33	83.51	84.07
Naïve Bayes	50.0	63.41	65.03	49.66
AdaBoost	92.4	91.13	88.67	89.71
Decision Tree	92.2	89.32	89.32	88.74
SVM	84.1	91.47	64.82	80.29
Random Forest	94.1	93.45	89.13	91.24
MLP	84.3	86.72	66.67	76.06
Deep Neural Network	94.2	90.08	69.37	74.19

Precision and harmonic mean recall Precision and harmonic mean F1-score.



Fig 2



Fig 3

According to the experimental results, the accuracy of the classification of ensemble based classifier such as the one provided by the Random Forest and AdaBoost is greater than the two individual ones. The high level of decision trees contributed to the fact that Random Forest exhibited high generalization. The AdaBoost reduced the biasness of the samples that were incorrectly classified; this was carried out through a reweighting regimen in a sequence of iterative processes to enhance the classification ability.

The Deep Neural Network model also achieved competitive outcomes because it learnt the non-linear correlation between the features of the transactions. The discrepancy between the recollection values of the various models was brought about by the imbalance in the data.




Fig 4

Such classifiers as Naive Bayes were not good to deal with because they used the features to assume that they are independent of each other. The suitable models were SVM and the Logistic Regression, which were not equally powerful as the projects with the ensemble models and deep learning models.

This assisted in providing a summary of the variations between the models in the magnitude of bar charts and performance tables with the assist of the graphical comparison. Confusion matrix development was also used to obtain false positive and false negative.

Overall, the results demonstrate that in the cases when the enlisting techniques of ensemble learning and deep learning are adopted, the rate of identifying fraud is significantly increased and the false positive is reduced.

TABLE 3

Comparative Performance of Machine Learning Models

Model	Accuracy (%)	Precision	Recall	F1 Score
Logistic Regression	91.3	0.89	0.86	0.87
Decision Tree	92.8	0.90	0.88	0.89
Support Vector Machine	93.6	0.91	0.89	0.90
Random Forest	95.2	0.94	0.92	0.93
Naïve Bayes	90.1	0.87	0.85	0.86
Neural Network (MLP)	94.4	0.92	0.91	0.91

V. DISCUSSION

This system uses the Decision Tree classifier that assigns banking transactions as either legitimate or fraudulent by constructing a hierarchical tree structure which consists of the split of features. It picks the most significant attribute of transaction at a node based on impurity metrics like the Gini Index and splits the data into smaller batches recursively until one reaches an ultimate classification at the leaf nodes. According to behavioural, frequency, and amount attributes of the transactions, the model is effective in their learning of the decision rules as it effectively captures the non-linear relationship in the data. Even though the Decision Trees are simple to interpret and computationally valid, they can suffer overfitting on an imbalanced dataset, so ensemble techniques are also believed to be used to enhance the overall performance and stability.

Discussion: Interpretation

Interpretation of the decision is the concept of knowing how the decision Tree model has classified a certain transaction. As the model is built in tree form, of sequential splits of feature description, every prediction can be followed through the tree of nodes to a leaf node by specification of decision rules. When studying the route followed in the tree, one can determine what the characteristics of the transaction, which include the sum, frequency, or behavioural trends, caused the ultimate decision that the transaction was either a fraud or a legitimate one.

This clarity causes the Decision Tree to be straightforward and useful in banking systems because financial institutions can have an easy time explaining why a transaction became cause of suspicion

VI . CONCLUSION

Conclusion

The current paper introduced a smart ensemble-based fraud detect system on banking transactions that is developed based on machine learning and deep learning. The researchers presented a main research goal that was to develop a reliable and scalable framework that will be able to detect fraudulent transactions on highly imbalanced banking datasets. The suggested system warrants a healthy comparison and meaningful performance analysis by delivering a well organized preprocessing chain, applying connected features normalization, and testing numerous supervised learning frameworks under the same setting.

The results of the experiment indicate that the traditional single classifiers (Logistic Regression and Naive Bayes) display moderate results, whereas the ensemble classifiers (Random Forest and AdaBoost) are much superior in terms of classification accuracy and stability. The fact that Deep Neural Network has been integrated also contributes to the ability of the system to identify complex and non-linear transaction patterns, which cannot be easily identified using traditional algorithms. The accuracy as well as the use of the measures of evaluation which include precision, recall, and F1-score are necessary to ensure that the model performance is evaluated comprehensively especially in the case of class imbalance where numbers of cases of fraud are few.

The results confirm that various learning paradigms used together would result in a higher level of fraud detection reliability and minimize false positive rates. This is particularly relevant when working in a banking setting, where the high level of false alerts can have a negative effect on customer experience and operational performance. The suggested framework is a compromise between the two extremes of security and practicality.

On the whole, it is possible to note that the smart ensemble-based design, generated within the framework of the current work, has high prospects of being implemented in the real world of contemporary digital banking systems. The system can help reduce financial losses, provide greater security of transactions and customer confidence in the online financial services by improving detection accuracy and making the model more robust.

VII . FUTURE WORK

Future Work: Research Directions

These areas of interest can be directed to future work, which can focus on enhancing the efficacy, flexibility, and scalability of the suggested intelligent fraud detection framework. Among improvements, one significant one is enhanced imbalance-handling policies, as the actual banking data is terribly biased toward legitimate business. To enhance the recall of fraud cases with low controlled false positive rates, techniques like SMOTE, adaptive synthetic sampling, cost-sensitive learning, and

ensemble resampling can be combined to enhance the recall of the fraud cases. It can be further optimized with systematic hyperparameter tuning with either Grid Search or Random Search or Bayesian Optimization to find the best parameter settings of each algorithm to increase the stability of the model and predictive performance across datasets. The second direction worth being taken is that of real-time detection of frauds systems which are able to stream the financial transactions. Combining streaming platforms and implementing the models on distributed cloud systems, the framework can be extended to support large amounts of transactions in minimal latency. This would render the system appropriate to the real-life banking practices where immediate fraud detection is essential to avoid loss of finances. Moreover, retraining mechanisms of the models can also be implemented, providing the ability to learn continuously based on the available new transaction data so that the ability to adapt to changing patterns of fraud and the appearance of new strategies of attackers is reached.

The further improvements can also focus on more advanced deep learning models like Long Short-Term Memory (LSTM) networks to get the sequence of transactions over time since fraud can have a temporal pattern. It is possible to research graph-based neural networks and examine their connections between accounts, devices, and the networks of transactions to identify organized or network-level fraud schemes. In addition, the interpretability can be enhanced by adding Explainable Artificial Intelligence techniques to enable financial institutions to learn why a transaction is considered a red flag.

Such transparency is needed to comply with regulations and establish trust with the customers. With such innovations, the proposed system may develop into an all-encompassing, scalable, and intelligent system to detect fraud that will fit the current digital banking ecosystem.

VII . REFERENCES

- [1] E. Ileberi, Y. Sun, and Z. Wang, "A machine learning based credit card fraud detection using the GA algorithm for feature selection," *Journal of Big Data*, vol. 9, no. 24, pp. 1–17, 2022.
- [2] I. Y. Hafez, A. Y. Hafez, A. Saleh, and A. A. Abd El-Mageed, "A systematic review of AI-enhanced techniques in credit card fraud detection," *Journal of Big Data*, vol. 12, 2025.
- [3] J. J. Assabil, "Credit card fraud detection using machine learning algorithms: A comparative study of six models," *International Journal of Intelligent Systems and Applications in Engineering*, vol. 12, no. 23, pp. 862–870, 2024.
- [4] P. K. R., R. Mathew, A. Walawalkar, P. Patil, U. Shirode, and A. Gaadhe, "A comparative analysis of machine learning techniques for detecting credit card fraud," *International Journal of Intelligent Systems and Applications in Engineering*, vol. 12, no. 3, pp. 146–153, 2024.
- [5] L. Zhang, "Credit card fraud detection based on machine learning algorithms," *Applied Science and Technology*, 2025.
- [6] Y. Lakshmi Durga, K. Samatha, S. Muskan, K. Jaya Shri, and A. Jahnavi, "Fraud detection in banking data by machine learning techniques,"

International Journal of Engineering Research and Science & Technology,
vol. 21, no. 2, pp. 1478–1481, 2025.

- [7] E. Ahmad, A. Bajpai, B. Venugopal, K. T. V. Rao, D. E., and R. N. Murthy, "AI-powered fraud detection and prevention in banking using deep learning," *Journal of Informatics Education and Research*, vol. 5, no. 2, 2025.
- [8] D. Xiao et al., "Revolutionizing fraud detection using machine learning in credit card security," *International Journal of Computers and Applications*, 2025