A Multi-Stage Hybrid Gradient Boosting Model With Optimal Selection Strategy For Online Fraud Detection

Lakhan Singh[1], Laxman Singh[2]
[1,2]*JB Institue of technology, Dehradun*

Abstract*—The increasing popularity of digital payment systems has led to a rise in online fraud, increasing the demand for accurate and scalable fraud detection solutions. Fraud detection remains challenging due to severe class imbalance between fraudulent and legitimate transactions, the continuously evolving nature of fraud patterns, and the requirement for transaction verification within milliseconds. A hybrid machine learning framework is proposed to detect fraudulent transactions by integrating XGBoost and CatBoost classifiers using a soft-voting ensemble strategy. Principal Component Analysis is applied to reduce input dimensionality while preserving 95% of the original data variance and enhancing model performance. XGBoost demonstrates strong learning capability on transformed numerical features, while CatBoost improves robustness through effective regularization and reduced overfitting. The hybrid framework combines probabilistic outputs from both classifiers to improve prediction confidence. Performance evaluation is conducted against individual boosting models using accuracy, precision, recall, F1-score, and ROC-AUC metrics. Experimental results indicate that the hybrid model achieves superior performance, including higher fraud recall and reduced false positives, demonstrating suitability for real-time online transaction fraud detection systems.

Index Terms*—Online payment fraud detection, machine learning, ensemble learning, XGBoost, CatBoost, Principal Component Analysis, class imbalance, real-time transaction analysis.

I. INTRODUCTION

With the rapid growth of digital payment systems has changed the way that financial transactions are completed by providing increased levels of efficiency, convenience, and accessibility. However, increased use of digital payment systems has also created a higher level of online payment fraud activity and therefore greater financial loss and security issues associated with online payment usage. Traditional fraud detection methods are predominantly rule-based and limited by predefined patterns and manual verification; as such, traditional methods are not well suited for large, high-dimension and highly imbalanced transactional databases, particularly when coupled with the emergence of new fraud strategies. Use of machine learning-based techniques for fraud detection can assist in providing a more effective alternative as these techniques provide the benefit of enabling the system to automatically detect complex patterns within historical transactional data, thereby providing improved accuracy concerning the detection of fraud over time. Gradient boosting ensemble learning has been effective in dealing with complex nonlinear relationships, categorical variables, and highly unbalanced datasets, as evidenced by many previous studies on its use. CatBoost and XGBoost have been widely utilized in the classification of financial data for their robustness, scalability and prediction capability and as such, they are well-suited for the fraud detection applications presented within this project. Additionally, the project creates an online fraud detection application for online payments that implements both CatBoost and XGBoost, as well as advanced preprocessing techniques, feature engineering and dimensionality reduction through Principal Component Analysis, resulting in computational efficiency while maintaining significant amounts of information.

II. LITERATURE REVIEW

[1] Dhiman et al. suggested a machine learning method for identifying fraudulent financial activity in real-time via the development and analysis of transaction patterns and features extracted from them. Experimental results indicate that the use of machine learning to detect fraud is effective; therefore, real-time monitoring systems must be utilized for

identifying fraudulent activities on digital payment platforms.

[2] Almazroi and Ayub developed a fraud detection model based on the use of multiple machine learning algorithms and evaluated the performance metrics associated with the model. The study used actual transaction data from real world transactions to evaluate the efficacy of the fraud detection system.The results indicate that the ensemble and boosting based algorithms outperformed other algorithms with regards to fraud detection, thus establishing the ability of machine learning to help reduce financial losses associated with online payment fraud.

[3] The research presented by Xiao et al. involves the performance comparison of CatBoost, XGBoost, and LightGBM for identifying fraudulent transactions. The study evaluated the strengths and weaknesses of each algorithm on fraud datasets. The study concludes that gradient boosting can effectively be used to identify fraudulent transactions because they are well suited to handling imbalanced data and complex feature interactions, therefore making them appropriate for large-scale fraud detection efforts.

[4] Babu et al. showed how they used Streamlit to build web apps that implement both XGBoost and CatBoost machine learning algorithms for the purpose of detecting credit card fraud. The authors found that CatBoost performs better with respect to categorical features, whereas both algorithms yield similar levels of performance.

[5] AlSagri proposed a hybrid multi-stage machine learning detection framework for identifying credit card anomalies and fraud. The experimental results showed that the proposed hybrid multi-stage machine learning framework produced significantly improved performance in the detection of credit card fraud when compared to individual detection models, making it an appropriate detection framework for the high-risk area of financial services.

[6] Malik et al. presented a hybrid machine learning architecture to detect credit card fraud. They proposed that their hybrid architecture improves the performance of fraud detection by analyzing both transaction patterns over an extensive period and transaction patterns in the immediate past. The results of theexperiments indicated that hybrid ensemble models outperform individual classifiers and provide greater robustness against extremely imbalanced datasets of fraud.

[7] The authors, Ileberi and Sun, recommended an ensemble approach based on deep learning to detect fraudulent use of credit cards. They utilized multiple architectures of deep learning to enhance their ability to classify accurately. Financial data sets are often very complex and large in size; the use of an ensemble-approach based on deep learning has been shown to significantly improve the accuracy and generalization performance of fraud detection, particularly in transactions of this type.

[8] Karthik et al., used hybrid ensemble modeling to capture behavioral patterns of customers when detecting fraudulent use of credit cards. By analyzing the pattern of transactions made over time, they identified anomalous behaviors during that period. The results of their study demonstrated that learning behavioral patterns of customers is extremely useful in improving the accuracy of detecting fraud in the financial realm.

[9] Khare and Srivastava published a thorough analysis of AI-based fraud-detection techniques that are being utilized for online transactions. In addition to outlining the various machine-learning methods that are available for detecting fraud, this study summarized the many ways these methods can be applied in fraud-prevention systems. Finally, the study highlighted the increasing use of AI-based technologies to improve the security of financial transactions conducted on the internet.

[10] Hashemi and colleagues studied how to use machine learning to identify fraudulent activities within banking transactions using a variety of classifier algorithms on actual transaction records from several banks. They assessed the utility of each algorithm compared to an established standard for detecting unusual patterns in bank transactions. Hashemi et al. found that when applying a machine learning framework for fraud detection, overall detection accuracy was significantly improved over previous methods, regardless of the type of algorithm used.

III. SYSTEM ARCHITECTURE

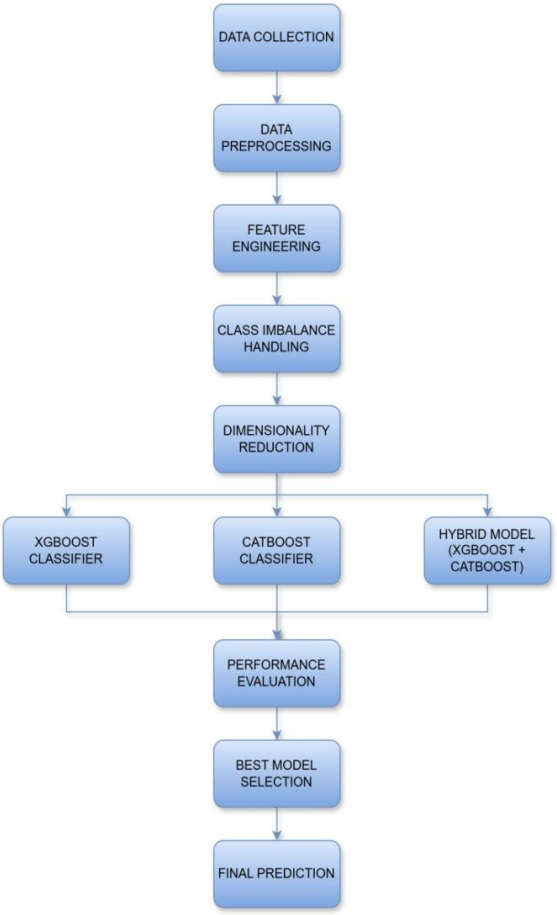

IV. PROPOSED METHODOLOGY

A. Data Source Unit

The Dataset derived from Kaggle for this research was a simulation of online payment transactions that is widely used for research about fraudulent activity detection and contains a vast quantity of data. All records contain both numeric and category data. The collection includes attributes such as transaction step, transaction type, amount, sender ID, receiver ID, and balances (before & after transaction) of each account. Among all the attributes in the Dataset, there are two Target Variables: isFraud to indicate whether the transaction was determined to be fraudulent or not and isFlaggedFraud to specify whether the transaction should be flagged as suspicious, based on predetermined criteria. The Dataset has an inherent Class Imbalance, only a very small number of records are associated with fraudulent activityand mimics the volume of transactions seen in legitimate online payments.

B. Data Preprocessing

Preparation was completed to provide a higher level of consistency in the dataset, therefore allowing it to produce higher quality results from the models. The dataset was evaluated to locate missing values, duplicated records, and possible inconsistencies in balances. Appropriate formats were established for categorical attributes, and numerical attributes were all treated with the same scaling, through normalization. Due to the very large disparity between the number of legitimate transactions and the number of fraudulent transactions, effective resampling methods were utilized to mitigate the problem. All features reflecting changes between before and after transaction balance were also evaluated, to determine the presence of abnormal behavior. Principal Component Analysis was used to minimize the number of dimensions, while retaining the maximum amount of variance, thus providing increased efficiency and generalization for the model.

C. Proposed Framework

The proposed methodology incorporates both XGBoost and CatBoost classifier methods, as well as a hybrid gradient boosting classifier architecturewhich also includes soft-voting.The resultant classifier will effectively identify fraudulent online transactions through independent training of three distinct classifiers on the transformed transaction features using PCA. Post the completion of the individual classifier training, the predictions from all three classifiers were merged to form a more effective and reliable approach to identifying fraudulent transactions. The resulting hybrid-gradient boosting architecture captures complex non-linear transaction behavior while reducing the effects of class imbalance commonly found in financial fraud detection data.

XGBoost Classifier

TheExtreme Gradient Boosting makes an excellent candidate to be a powerful non-linear pattern learner for modeling the complex interaction of numerous attributes of a given transaction. XGBoost employs the principles of additive ensemble methods; that is, each decision tree built on top of the last will correct the mistakes made by the trees prior to it. Given a

transaction's feature vector x_i, the predicted value of that transaction is as follows:

$$\hat{y}_i = \sum_{m=1}^{M} f_m(x_i) \qquad (1)$$

here each $f_m(\cdot)$represents a regression tree. The model is trained by minimizing a regularized objective function:

$$\mathcal{L}^{(t)} = \sum_{i=1}^{N} l(y_i, \hat{y}_i^{(t-1)} + f_t(x_i)) + \Omega(f_t) \qquad (2)$$

Here, $l(\cdot)$ denotes the loss function measuring prediction error, and $\Omega(\cdot)$is a regularization term that controls model complexity. The regularization is defined as

$$\Omega(f) = \gamma T + \frac{1}{2}\lambda \sum_{j=1}^{T} w_j^2 \qquad (3)$$

The equation specifies that *T* is the number of leaf nodes, wj the weight of each leaf node. This way, XGBoost can effectively learn arbitrary non-linear fraud behavior including anomalous transaction amounts, balance fluctuations, etc., while preventing overfitting.An additional feature of XGBoost is that it uses an iterative process of error correction so that detection can focus on hard-to-classify fraud and improve performance as a result.

CatBoost Classifier

Fraud detection is enhanced and improved using a CatBoost Classifier to improve both generalization as well as robustness. CatBoost utilizes the technique called ordered boosting with the use of symmetric (oblivious) decision trees to help reduce the prediction bias and variance during the training process through the addition of a gradient boosting algorithm. As with XGBoost, the predicted results can be expressed in an additive model:

$$\hat{y}_i = \sum_{m=1}^{M} f_m(x_i) (4)$$

To avoid biased gradient estimation, CatBoost computes predictions at iteration t using only previously trained trees:

$$\hat{y}_i^{(t)} = \sum_{k<t} f_k(x_i) (5)$$

Through ordered boosting, you can achieve stability in the learning phase and enhance generalization when faced with a heavily imbalanced fraud data set. By applying advanced regularization techniques, CatBoost helps to identify less obvious fraud patterns that may not be detected by other classification models. Due to this, CatBoost performs exceptionally well at recognising rare instances of fraud while maintaining low false positive rates.

Hybrid Soft-Voting Ensemble

A soft-voting strategy is used to create a hybrid ensemble model by using a combination of both XGBoost and CatBoost models to provide higher sensitivity and specificity in detecting fraud. Rather than just using one classification model to make a prediction, the ensemble model will provide an aggregate output of both classifiers' probability outputs. Therefore, the final estimate of an instance being fraudulent will be based upon:

$$P_{ensemble}(y_i = 1) = \frac{1}{2}\big(P_{XGB}(y_i = 1) + P_{Cat}(y_i = 1)\big) \qquad (6)$$

The final class label is determined by applying a threshold
τ:

$$\hat{y}_i = \begin{cases} 1, & P_{ensemble} \geq \tau \\ 0, & \text{otherwise} \end{cases} (7)$$

The ensemble strategy combines the complementary learning methodologies of XGBoost and CatBoost, as well as decreases model-specific biases and increases the global detection accuracy. The ensemble model achieves greater fraud recall and reduced false positive rates by leveraging the models' probability predictions, making it ideal for real-time fraud detection in online transactions.

Performance-Based Model Selection

The evaluation of XGBoost, CatBoost and the hybrid ensemble, which occur after training, is done by using classification metrics (precision, recall, F1 score, and ROC – AUC). The final decision model is chosen based on the model with the highest detection effectiveness, particularly in terms of recall and ROC – AUC for fraud detection. The adaptive selection method allows for better performance as transaction patterns change and fraudulent behavior evolves.

V. EXPERIMENTAL SETUP

A. Feature Configuration and PCA Setup

In developing a fraud detection system, eight transaction input features will be analyzed; these features are the transaction step, transaction type, transaction amount, and account balance attributes associated with the transaction. To eliminate the complexity of calculating the numerous features and to eliminate duplicated features, Principal Component Analysis has been applied to the normalized feature space. From the PCA model that used the normalized feature space input, the number of dimensions was reduced to three principal components while still retaining approximately 95% of the variance from the original feature space. This process has allowed for increasing the efficiency of the models and improving the model's generalization performance without significant loss of information.

B. Training and Evaluation Settings

A large-scale online payment transaction dataset consisting of 1,048,576 records with highly imbalanced fraud distribution was used as part of the experimental evaluation. The dataset was divided into two separate subsets, one for the purpose of training and one for testing the different models. An 80:20 ratio specified how to divide the data between training and testing subsets. Training on the test data was done the same way for XGBoost, CatBoost, and the proposed hybrid ensemble model - all were trained in similar experimental conditions with hyperparameters optimized to make their performance comparable. A standard set of evaluation metrics includes ROC–AUC, precision, recall, and F1-score was used to evaluate the performance of all models. All experiments were conducted using Python-based machine learning libraries.

VI. RESULT AND DISCUSSION

A. Real-Time and Batch Inference Performance

Figure 5.1 shows how the proposed system implements a single transaction prediction. Users must enter specific individual transaction attributes liketransaction step, type, amount, and account balances into the system, which then classifies these inputs using three different machine learning algorithms CatBoost + PCA, XGBoost + PCA, or an Ensemble Model that combines both classifiers. This mode allows users to receive an immediate fraud assessment of an individual transaction while continuing to use this technology within the ecosystem of an online payment platform.

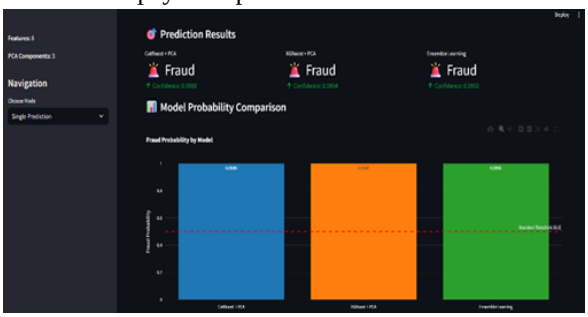

Figure 5.1 Single Transaction Prediction Analysis

The proposed solution will ensure that all users utilize the same format of features and receive accurate predictions.

A batch prediction operation was performed on a dataset containing online payment transactions to analyze the large-scale performance. The prediction results table contains transaction details along with the fraud labels and predicted probabilities. Consistently predicting probabilities across all transactions indicates that the model output is stable. The use of PCA helps reduce inference time, which indicates that the predicted model is appropriate for both real time and high-volume applications in the finance industry.

B. Fraud Detection Performance Comparison

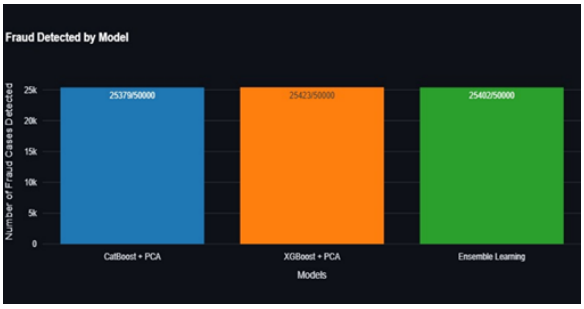

Figure 5.2Fraud detection count comparison.

As illustrated in Fig 5.2, a comparison of fraud case detection across multiple model types has been made. Total number of fraud cases detected by each model indicates that CatBoost + PCA detected 25,379 total fraud cases, while XGBoost + PCA detected 25,423. Ensemble learning detected 25,402 total fraud cases, thus providing an equivalent level of performance but increased stability through the reduced degree of

model-specific biases. Overall, ensemble learning models improve both the overall reliability of fraud case detection and reduce the model-specific bias.

C. Fraud Probability Distribution Analysis

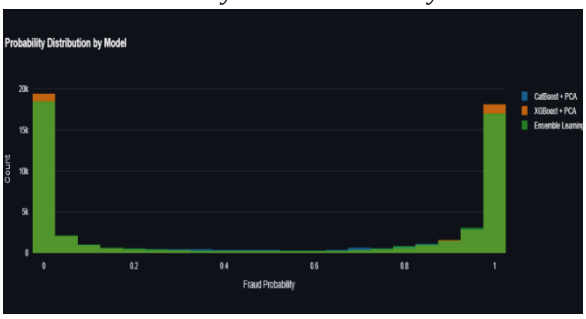

Figure 5.3 - Fraud Probability Distribution Analysis

In Figure 5.3, all models' probability distributions display a clean bimodal pattern where prediction probabilities are concentrated near both ends 0s & 1s. This implies that these models have an exceptional degree of confidence in their classification decision-making abilities. In addition, the ensemble model's probability distribution develops a smoother transition between the low and high probability zones than the other models, with far less uncertainty around the middle of each probability range. As a result, the ensemble model's outputs are more accurately calibrated and therefore represent a more reliable source than any of the other models. The separation of positive and negative predictions via probability is critical for minimizing the incidence of false positives in fraud detection applications.

D. Model Agreement Analysis

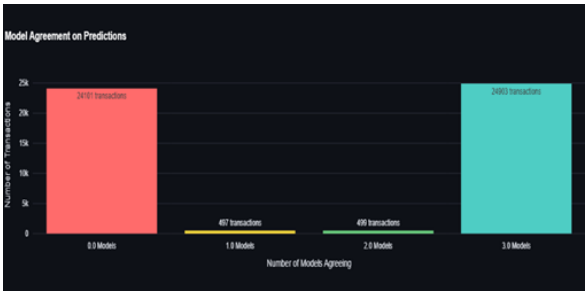

Fig. 5.4. Model prediction agreement distribution.

Figure 5.4 illustrates how well XGBoost, CatBoost, and the hybrid ensemble areproposing agree in their predictions. It shows that most of the transactions have perfect agreement from all three models and therefore these transactions have extremely high classification reliability and stable decision-making behavior. Only a small number of transactions show some degree of agreement, which reflects uncertainty or borderline transaction patterns. This high level of agreement between models demonstrates that the ensemble strategy has a significant impact in removing model-specific bias and increasing the overall robustness of all prediction models when used for real-time fraud detection application purposes.

E. Computational Complexity Analysis of Fraud Detection Models

The analysis highlights the scalability and computational efficiency of tree-based boosting models for large-scale fraud detection.

Table 5.1 - Model Complexity

MODEL	AUC	PREC.	REC.	F1.
CatBoost	0.985	0.92	0.89	0.905
XGBoost	0.982	0.91	0.87	0.89
Hybrid	0.988	0.93	0.9	0.915

Here, N denotes the number of transactions, k represents the number of features after PCA transformation, M is the number of decision trees, M_X and M_C denote the number of trees in XGBoost and CatBoost respectively, and d represents the maximum depth of the trees.

Performance Evaluation and Discussion

To evaluate the effectiveness of a hybrid gradient boosting system in detecting online fraud, various standard classifications were utilized. Figure 5.5 summarizes the comparative performance of XGBoost, CatBoost, and the proposed hybrid ensemble method.

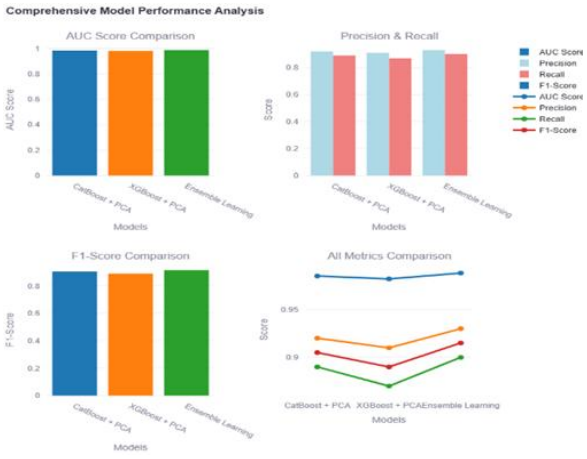

Figure 5.5 Model Performance Comparison

The experimental results were obtained by evaluating all models on a held-out test dataset after training with optimized hyperparameters and dimensionality reduction using PCA.

Table 5.2 – Performance evaluation

MODEL	TIME	SPACE	SCALE
XGBoost	O (MNk logN)	O(Md)	High
CatBoost	$O(MNk)$	O(Md)	High
Hybrid	O (M_XNk logN+M_CNk)	O(M_X +M_C) ·d)	High

It was noted that the hybrid ensemble model produced the best results over the various evaluation metrics; this model had an AUC of 0.988, a precision score of 0.93, a recall score of 0.90, and an F1 score of 0.915. The increase in performance indicates the successful application of using the soft voting method to combine XGBoost and CatBoost into one ensemble model. Although both CatBoost + PCA and XGBoost + PCA performed well individually with CatBoost + PCA having a slight performance advantage, the hybrid model achieved significantly higher recall and F1 score values, which indicates that it is more effective at detecting fraudulent transactions while reducing the chances of false positives. The ensemble model's increased recall value shows that it can better identify minority instances of fraud, which is an important consideration when working with highly imbalanced financial datasets.

VII. CONCLUSION

The research proposes a bimodal method of creating a fraud detection algorithm to detect fraudulent online payments. The bimodal method uses three classifiers (XGBoost, CatBoost and PCA) to generate an Ensemble Classifier that produces improved performance when compared to the other three classifiers and their individual abilities to detect and prevent fraud occurring as part of an online transaction. The study also showed that an Ensemble Classifier combined with boosting algorithms and feature reduction through PCA provided an efficient means of detecting and filtering a large volume of online transactions during the event. Future research will focus on how to incorporate streaming data into the hybrid model developed in this study, as well as how to integrate Explainable AI into the hybrid model to improve its adaptability and transparency in fraud detection.

REFERENCES

[1] D. Dhiman, A. Bisht, A. Kumari, H. Anandaram, S. Saxena, and K. Joshi, "Online fraud detection using machine learning," in *Proceedings of the International Conference on Artificial Intelligence and Smart Communication (AISC)*, 2023, pp. 161–164.

[2] A. A. Almazroi and N. Ayub, "Online payment fraud detection model using machine learning techniques," IEEE Access, vol. 11, pp. 137188–137203, 2023.

[3] Y. Xiao, L. Tan, and J. Liu, "Application of machine learning model in fraud identification: A comparative study of CatBoost, XGBoost and LightGBM," Preprints, 2025.

[4] E. Babu, J. J. Maliakal, D. Babu, et al., "Performance comparison of XGBoost and CatBoost algorithm in credit card fraud detection with Streamlit-based web application," in Proceedings of the International Conference on Advancement in Renewable Energy and Intelligent Systems (AREIS), 2024, pp. 1–6.

[5] H. S. AlSagri, "Hybrid machine learning based multi-stage framework for detection of credit card anomalies and fraud," IEEE Access, 2025.

[6] E. F. Malik, K. W. Khaw, B. Belaton, W. P. Wong, and X. Chew, "Credit card fraud detection using a new hybrid machine learning architecture," Mathematics, vol. 10, no. 9, Art. no. 1480, 2022.

[7] E. Ileberi and Y. Sun, "A hybrid deep learning ensemble model for credit card fraud detection," IEEE Access, 2024.

[8] V. S. S. Karthik, A. Mishra, and U. S. Reddy, "Credit card fraud detection by modelling behaviour pattern using hybrid ensemble model," Arabian Journal for Science and Engineering, vol. 47, no. 2, pp. 1987–1997, 2022.

[9] P. Khare and S. Srivastava, "AI-powered fraud prevention: A comprehensive analysis of machine learning applications in online transactions," Journal of Emerging Technologies and Innovative Research, vol. 10, no. 9, pp. 2349–5162, 2023.

[10] S. K. Hashemi, S. L. Mirtaheri, and S. Greco, "Fraud detection in banking data by machine learning techniques," IEEE Access, vol. 11, pp. 3034–3043, 2022.