

Online Payment Fraud Detection Using Machine Learning

Mr. Pushparaj P¹, Pravin Rahul S K², Navedh Akhtar Jamali N³, Ranjithkumar S⁴

¹Assistant Professor/Department of Artificial Intelligence and Data Science Erode Sengunthar Engineering College, Erode, India

^{2,3,4}UG Scholar/Department of Artificial Intelligence and Data Science Erode Sengunthar Engineering College, Erode, India

DOI: <https://doi.org/10.51583/IJLTEMAS.2026.150400082>

Received: 10 April 2026; Accepted: 15 April 2026; Published: 12 May 2026

ABSTRACT

Online payment fraud detection is a critical problem in the financial sector due to the increasing volume of digital transactions. This paper proposes a machine learning-based fraud detection system using CatBoost, XGBoost, and a soft voting ensemble model. Principal Component Analysis (PCA) is applied for dimensionality reduction, and SMOTE is used to address class imbalance. The models are evaluated using precision, recall, F1score, and AUC. Experimental results show that the ensemble model outperforms individual models with improved accuracy and robustness. A real-time fraud detection system is also developed using Streamlit to support both single and batch predictions. The proposed system demonstrates high efficiency and scalability for practical applications.

Keywords: Fraud Detection, Machine Learning, CatBoost, XGBoost, PCA, SMOTE, Ensemble Learning

INTRODUCTION

Credit cards have transformed how we make purchases and manage our finances. A Online payment is a plastic payment card provided by a financial organization, usually a bank that allows the cardholder to borrow money to purchase goods and services. Unlike debit cards, which remove cash straight from the cardholder's bank account, credit cards provide a line of credit that must be returned at a later date, typically with interest. The cardholder may use the Online payment for a variety of purchases, both in person and online, making it a convenient and generally recognized payment method across the world [1] [2]. Credit cards have become an essential aspect of modern consumer culture, giving people more purchasing power and the freedom to manage their finances properly. When a consumer applies for a Online payment and is authorized, the financial institution sets them a credit limit, which is the most money they may borrow with the card. The cardholder may use the Online payment to make purchases up to the limit. Each month, the Online payment firm delivers a billing statement outlining the transactions that occurred during that time period as well as the minimum amount owing. [4] [3] While paying the minimal amount keeps the account in good standing, it is best to pay off the whole debt to prevent interest costs. If the cardholder fails to make the minimum payment on time, they risk incurring late penalties and harming their credit score. Credit cards sometimes come with a variety of incentives and advantages, such as cash back, travel points, or discounts, to encourage card use and loyalty [7]. Fraud detection is a vital use of data analysis and machine learning that aims to identify and prevent fraudulent activity and transactions across several domains. As technology progresses and financial transactions move to digital platforms, the danger of fraudulent activity increases, making fraud detection more important than ever [8].

As the finance sector continues to undergo digital transformation, online payment systems are now ubiquitous in modern business, giving users unprecedented speed and convenience. But the growth of this convenience has simultaneously provided the perfect conditions for sophisticated fraudulent behaviour which causes a critical risk that can result in significant global financial losses for individuals and organizations. The development of Online Payment Fraud Detection Systems has consequently become a top priority to protect the integrity and security of the digital economy. Unlike traditional, inflexible rulebased solutions that can't keep up with the rapid

pace of change of techniques by cybercriminals, contemporary solutions such as this one use contemporary machine learning (ML) algorithms to analyse complex transactional data. [6]ML models are capable to detect subtle patterns and anomalies associated with fraudulent behaviours in real time, mitigating risk to the reliability of transactions and reducing false positives.

Whether its Online payment fraud, insurance fraud, identity theft, or internet scams, companies and financial institutions use sophisticated fraud detection systems to protect their assets, consumers, and preserve faith in their services. Fraud detection systems examine massive volumes of transactional and behavioural data to find anomalies and suspect trends. Fraud may have serious consequences, including financial losses, tarnished reputations, and impaired consumer trust.

Fraud can have serious consequences, including financial losses, damaged reputations, and compromised customer trust. Fraud detection systems play an important role in avoiding such repercussions by proactively identifying and flagging questionable transactions. Effective fraud detection benefits financial organizations by not just protecting their assets but also ensuring regulatory compliance and increasing client confidence.

Related Work

Maryam Habibpour [1] and colleagues propose in this work Numerous research have used deep neural networks (DNNs) to identify Online payment fraud, with the goal of improving point prediction accuracy and avoiding unwanted biases through the development of various network architectures or learning models. It is critical to measure uncertainty in conjunction with point estimate because it lowers model unfairness and allows practitioners to build dependable systems that prevent making incorrect judgments due to uncertainty. Because fraudsters continuously change their strategies, DNNs meet observations that do not come from the same process as the training distribution. Furthermore, because to the time-consuming nature of the procedure, only few transactions are reviewed by experienced specialists in order to update DNNs. These characteristics make it necessary to clearly evaluate the uncertainty associated with DNN predictions in real-world card fraud detection scenarios.

In this paper, Asma Cherif [2] et al. Online payment fraud is becoming a serious and growing problem as new technologies and communication channels emerge, such as contactless payment. This article gives a comprehensive examination of current research on detecting and forecasting fraudulent Online payment transactions from 2015 to 2021. The 40 papers picked for consideration are analysed and classified based on the topics they cover (class imbalance problem, feature engineering, etc.) and the machine learning approach they use (conventional and deep learning modelling). Our analysis reveals that deep learning has received little research, implying that more research is needed to address the difficulties in detecting Online payment fraud using cutting-edge technologies such as big data analytics, large-scale machine learning, and cloud computing. Our study is a significant resource for academic and industrial researchers in analysing financial fraud detection systems and designing trustworthy solutions by highlighting existing research challenges and future research opportunities.

Dr. Tran Khanh Dang [3], ET. The issue of unbalanced datasets is a fundamental concern for constructing reliable Online payment fraud (CCF) detection systems, as stated in this system. In this paper, we study and evaluate current advances in deep reinforcement learning (DRL) and machine learning (ML) algorithms for CCF detection systems, including fraud and non-fraud labels. The imbalanced CCF dataset is resampled with SMOTE and ADASYN, two resampling methods. This balanced dataset is then exposed to ML algorithms to generate CCF detection models. The imbalanced CCF dataset is then used to build detection algorithms with DRL.

Jacobo Chaquet-Ulldemolins [4] et al. proposed this system. Artificial intelligence (AI) has lately gained popularity in the global economy due to its exceptional ability to analyse and model data in a variety of disciplines. As a result of this condition, society is rapidly becoming more automated, and these new approaches may be combined to form a beneficial tool for addressing the difficult challenge of credit fraud detection. However, severe restrictions make it impossible for financial institutions to comply with them while employing modern procedures. From a methodological approach, auto encoders have demonstrated effectiveness in finding

nonlinear features in a range of problem domains. However, auto encoders are opaque and often referred to as "black boxes." In this study, we provide an interpretable and impartial CFD approach.

Esra Faisal Malik [5], for example. As mentioned in this article financial crimes have progressively harmed financial institutions. Various single and hybrid machine learning algorithms have been used to detect crimes such as Online payment fraud. However, due to a lack of additional research on alternative hybrid algorithms for a specific dataset, these techniques have significant limitations. This paper proposes and tests seven hybrid machine learning models for detecting fraudulent acts on a real-world dataset. Modern machine learning techniques were initially applied to detect Online payment fraud, and the best single algorithm from the first phase was used to create the hybrid approaches. The hybrid models created were separated into two phases. Our results revealed that the hybrid model Adaboost + LGBM is the best model due to its superior performance. Future study should focus on exploring different hybridization strategies and Online payment domain algorithms.

In their article, Ibtissam Benchajiet [6] discusses a new system for Online payment fraud detection, making improvements to current testing methods through the use of sequential modelling in traditional machine learning. Because the effectiveness of any fraud detection technique depends on the features available for modelling, the challenges & limitations of transactional payment data must be explored. Moreover, the study highlights the importance of information or features, even in the form of time series data that represent transactions; the detection of fraud is highly contingent on the presence of some essential predictive characteristics or information. In addition, the model proposes the framework to be robust against fraudulent activities within a transactional dataset, allowing for techniques that can be developed to optimize values that may not be present or available within the dataset. The methodology for achieving this will be strengthened the combination of three costeffective probabilistic dimensionality reduction feature cross-validator selection (LSTM). As Online payments are increasingly common, it should not be surprising that fraudulent activities are occurring. To properly combat and address fraud, financial institutions must take steps to enhance their monitoring systems to decrease significant losses. The proposed model being summarized should attempt to address scam fraud activities.

In this paper, Ebenezer Esenogho [7] and colleagues have suggested, "Recent improvements in e-commerce and communication systems have played a major role in the increased usage of credit cards both for online purchases, as well as traditional shopping. Nonetheless, the rate of fraudulent activity in online payment transactions has steadily risen, leading to considerable losses for financial institutions worldwide. Creating efficient and reliable fraud detection algorithms is important to minimizing financial loss resulting from fraudulent online payment transactions; however, there are challenges, as most online payment datasets are characterized by severe class imbalance. Additionally, using traditional machine learning algorithms for online payment fraud detection is not efficient: machine learning algorithms described and applied in this manner are predetermined by their design, which involves a static mapping of input vector to output vector, and cannot accommodate for the dynamic purchasing habits of online payment customers.

According to Samaneh Sorournejad [8] et.al. in this article, Online payment plays a detrimental and significant role in today's economy, as it has become an unavoidable element of household, business, and global activities. Although credit cards can have tremendous advantages when used properly with restraint, significant damage to credit and finances can be created through fraud. Numerous methods have been proposed to resolve the growth of Online payment fraud, but all have the same purpose of detecting Online payment fraud. Each of these techniques has certain drawbacks, advantages, and features. The focus of this paper is to address the challenges of Online payment fraud detection and to review the state of the art in Online payment fraud detection techniques, and datasets, and evaluation criteria. The advancements and drawbacks of the fraud detection techniques will be presented and compared. In addition, a taxonomy of the aforementioned techniques is presented, distinguishing between two principal strategies for fraud detection, namely, misuses (supervised) and anomaly detection (unsupervised).

In this paper, Yue Tian [9] et.al. Have introduced several machine learning methods that provide efficient transaction fraud detection, which is crucial to both individual and bank financial security. However, many current methods use only original features or use manual feature engineering or both, which prevents them from

learning discriminative representations from transactions. Since fraudsters often commit fraud by mimicking cardholders' behavior, the performance of existing fraud detection models is poor. Here we present an Adaptive Sampling and Aggregation-based Graph Neural Network (ASA-GNN) that learns discriminative representations in order to optimize transaction fraud detection performance. A neighbour sampling strategy used to remove noisy nodes and supply information to fraudulent nodes.

In this study, Imane Sadgali [10] et.al. suggests that as banking transactions continue to evolve, so does the risk of online payment fraud, particularly with the rapid technological advances we are witnessing; fraudsters are becoming more sophisticated and are continuously finding ways to work around the preventive models employed by financial systems. Several studies have developed predictive models for Online Payment Fraud detection using different machine learning methods. In this work, we propose an adaptive framework that enhances Online Payment Fraud detection using models that have recently demonstrated high levels of accuracy and that integrate the type of transaction and the clients profile.

Existing Work

In today's digital economy, credit cards are indispensable, and as their use has recently increased significantly, so has Online payment theft. Algorithms for machine learning (ML) have been used to detect Online payment fraud. However, it has proven challenging for ML classifiers to function at their best due to Online payment holders' dynamic shopping habits and the issue of class imbalance. This paper presents a robust deep-learning method to address this issue, which combines a multilayer Perceptron (MLP) as the meta-learner with long short-term memory (LSTM) and gated recurrent unit (GRU) neural networks as base learners in a stacking ensemble architecture. To balance the class distribution in the dataset, the hybrid synthetic minority oversampling methodology and edited nearest neighbor (SMOTE-ENN) method are used. According to the experimental findings, the suggested deep learning ensemble in combination with the SMOTE-ENN method produced sensitivity and specificity values of 1.000 and 0.997, respectively, which are better than those of other commonly employed ML classifiers and techniques in the literature.

Proposed System

The Online Payments Fraud Detection System follows a methodical approach to reliably identify fraudulent transactions. First, transactional datasets are gathered and uploaded, which include features such as transaction type, amount, account balances, and target fraud label. The feature names can either be uploaded from a file, or taken from the dataset to maintain consistency. The data must also be cleaned for completion and correctness in order to be prepared for modeling. Principal Component Analysis (PCA) is used to mitigate computational complexity, while improving model performance. PCA will convert the original high-dimensional feature space into a lowerdimensional representation that retains 95% of the variance. This step assists in minimizing the potential redundancies and allows the models to focus on the more informative components that signify fraud patterns. In the predictive modeling section, three approaches to model fitting are utilized - CatBoost with PCA, XGBoost with PCA, and an Ensemble model which combines both by utilizing soft voting.

CatBoost addresses working with categorical features very well and is often reliable with imbalanced data, while XGBoost excels at finding complex patterns in data and scaling the predictions. The Ensemble model combines the advantages of both models, resulting in enhanced outcomes that are more trustworthy and more resilient. The models' performances are determined using AUC, Precision, Recall, and F1Score, which permit comparison and selection of the best model for deployment. Both the models, PCA transformer, and feature metadata are used to deploy a web application based on Streamlit supporting transactions' single and batch predictions. The application provides the probabilities of fraud, confidence scores, probability distributions, and which model agreed with each prediction. There is interactive visualizations included, such as bar charts, radar charts, and feature tables to aid with interpretation and allow the user to see how the attributes of the transaction affected the prediction. Overall, the approach provides a system for fraud detection for online payment platforms that is reliable, scalable, and user-friendly, while integrating data preprocessing, dimensionality reduction, advanced machine learning, and interactive visualization.

Description of the dataset

The dataset used in this study is a publicly available online payment fraud detection dataset obtained from Kaggle. It contains a large number of financial transactions with both legitimate and fraudulent activities.

The dataset consists of approximately 50,000 transactions, with a highly imbalanced class distribution where fraudulent transactions represent a very small percentage of the total data. This imbalance makes fraud detection a challenging problem.

Each transaction includes features such as transaction type, amount, origin and destination account balances, and time step. The target variable is “isFraud”, which indicates whether a transaction is fraudulent (1) or legitimate (0).

To ensure proper evaluation, the dataset was divided into training and testing sets using an 80:20 stratified split, preserving the class distribution in both sets.

Data pre processing

After the datasets are loaded, we perform preprocessing in order to prepare the data for the modeling process. This process includes handling any missing or irrelevant values, along with preserving the ordering of features for reliable input. We then proceed with a dimensionality-reduction technique known as Principal Component Analysis (PCA) to reduce the dataset dimensionality while preserving 95% of the variance. Implementing PCA will decrease computation costs and multi collinearity, ensuring a clean and optimal set of features for the models to learn the patterns of fraudulent transactions.

Model Selection and Training

Three predictive models are applied:

1. CatBoost + PCA: CatBoost will be used since it can handle categorical variables well and does not over fit on imbalanced data.
2. XGBoost + PCA: XGBoost captures complex relationships within the data and generates fast, high-quality predictions.
3. Ensemble Learning (Voting Classifier): Soft voting is used to combine the probabilistic outputs of CatBoost and XGBoost to gain reliability and robustness.

All models are trained on the PCA-transformed features, and the relevant hyper parameters were selected to create better classification performance. The ensemble model will use the two models combined to reduce errors and increase fraud detection.

Feature Name	Data Type	Description
step	Integer	Unit of time in the dataset (1 step = 1 hour)
type	Categorical	Type of transaction: CASH-IN, CASH-OUT, DEBIT, PAYMENT, TRANSFER
amount	Float	Amount of the transaction
oldbalanceOrg	Float	Initial balance of the origin account before the transaction
newbalanceOrig	Float	New balance of the origin account after the transaction
oldbalanceDest	Float	Initial balance of the destination account before the transaction
newbalanceDest	Float	New balance of the destination account after the transaction
isFraud	Integer	Target label indicating fraud (1) or legitimate (0)

Table 1. Dataset Statistics

Model Evaluation

Standard assessment metrics are used to evaluate model performance: AUC (Area under the Curve), Precision, Recall, and F1-Score, which assess the model's ability to identify fraudulent transactions correctly while minimizing false positives and false negatives. The comparative evaluation ensures that the selected model or ensemble is the best performing candidate for deployment.

Deployment in Streamlit

The trained models, as well as PCA transformers and feature metadata, are saved and deployed in a Streamlit-based web app. The app supports:

- Single Transaction prediction: Users provide feature values and receive a probability score and classification for fraud detection.
- Batch Prediction: Users upload CSV files with transaction records for batch processing.
- Interactive visualizations: Users see probability distributions, charts that demonstrate level of agreement between models, and radar charts to help interpret model predictions.
- Feature transparency: Feature descriptions and counts are displayed for interpretability.

Results Interpretation and Decision Support

The system predicts fraudulent transactions and provides confidence scores and model agreement analyses, allowing decision-makers to focus their investigations on the riskiest transactions. Predictions based on an ensemble of models increase confidence; visualizations provide operational clarity.




Figure 1. System Block Diagram




Figure 2. System flow Diagram

RESULT AND DISCUSSION

The Online Payments Fraud Detection System was tested with three methods: CatBoost + PCA; XGBoost + PCA; and an Ensemble that combined both methods use soft voting. The models were trained on a dataset with transactional features including transaction type, amount, and account balances to predict fraud (the target). PCA was applied to the features to reduce their space while retaining 95% of the variance, which assisted in reducing the computational load of the program and models. In search of CatBoost + PCA, the model's performance was good handling categorical features and an imbalanced target. The classification report demonstrated very high levels of precision and recall for both legit and fraud classes, and the AUC score was 0.985 suggesting that the models accurately distinguished money laundering from non-fraudulent transactions.

Although the AUC for the XGBoost + PCA model slightly lowered to 0.982, it still performed notably well - showing how XGBoost captured complex patterns in the data and provided predictions efficiently and quickly. While precision and recall were lower than the CatBoost model, XGBoost remained competitive in fraud detection. The

Ensemble model achieved the best overall performance based on the AUC score of 0.988, which combined CatBoost + XGBoost predictions through soft voting. An ensemble model utilizes the best abilities of both models and reduced individual weaknesses when producing predictions that were collectively stronger across the sets of predictions. Evaluation of probability, model agreement, and batch prediction, indicated that the ensemble benefited by performing consistently higher than both models across all evaluations. In the Streamlit app, interactive visualizations such as bar charts, radar charts, and probability distributions were available to support final model reporting, detection of behaviour in models, and prediction next to single transactions.

Additionally, a confidence score was displayed as an interpretation of single transactions predicting likelihood of fraud, while batch predictions calculated aggregate statistics and offered file download. Overall findings indicated that the proposed system, based on the evaluated models, was the best predictor model in terms of ability, scalability, and interpretation of fraud detection that could be utilized in online payment.

Parameters for Evaluation

Accuracy

Accuracy measures the proportion of correctly predicted online payment approvals and rejections out of the total applications processed by the model.

Formula:

$$\text{Accuracy} = (\text{TP} + \text{TN}) / (\text{TP} + \text{TN} + \text{FP} + \text{FN})$$

Where:

- TP (True Positives): Number of Online payment applications correctly predicted as approved.
- TN (True Negatives): Number of Online payment applications correctly predicted as rejected.
- FP (False Positives): Number of Online payment applications incorrectly predicted as approved, but were actually rejected.
- FN (False Negatives): Number of Online payment applications incorrectly predicted as rejected, but were actually approved.

Precision

Precision evaluates the accuracy of the model's positive predictions – i.e., how many of the applications predicted as approved were actually approved

Formula:

$$\text{Precision} = \text{TP} / (\text{TP} + \text{FP})$$

Recall

Recall measures the model's ability to correctly identify all truly approved applications from the dataset

Formula:

$$\text{Recall} = \text{TP} / (\text{TP} + \text{FN})$$

F1 Score

The F1 Score is the harmonic mean of precision and recall. It offers a single performance metric that accounts for both false approvals and false rejections, making it suitable when there's class imbalance. Formula:

$$\text{F1-Score} = 2 \times (\text{Precision} \times \text{Recall}) / (\text{Precision} + \text{Recall})$$

Model	Precision	Recall	F1-Score	AUC Score
CatBoost + PCA	0.92	0.89	0.905	0.985

XGBoost + PCA	0.91	0.87	0.890	0.982
Ensemble (CatBoost+XGB)	0.93	0.90	0.915	0.988

Table 2. Performance table




Figure 2. Performance comparison

Single Prediction

Through the Single Prediction module, the user can input feature values of a single transaction and receive fraud predictions from three models: CatBoost + PCA, XGBoost + PCA, and Ensemble Learning. For every transaction, the system adds:

A classification (Fraud or Legitimate) based on a 0.5 probability threshold. o Confidence scores that are the probability of fraud for each of the individual models, enumerating the strength of their prediction.

Overall, the CatBoost + PCA model provided accurate transaction classifications for many of the categorical features, providing high confidence for transactions flagged for fraud. The XGBoost + PCA detected more complex fraudulent patterns and provided additional independent support to the analysis, although it had differences in confidence values in certain instances where the CatBoost flagged transactions for fraud. The Ensemble model provided probability averages of predictions made by the CatBoost and XGBoost. It consistently provided the most reliable classification all with the highest confidence score and consistently minimized the misclassification probability. Probability bar chart comparisons provided a quick visual check on how each model was assessing the transaction, which supports interpretability.

Overall, for single transactions, the system provided high accuracy and interpretability by providing:

- o Fraudulent transactions above 0.85 confidence scores. o Legitimate transactions above 0.90 confidence scores.
- o Consistently better prediction stability and classification than the individual models.
- o Predicted probabilities of fraud and classified labels from both CatBoost and XGBoost, as well as the Ensemble model.

- Aggregate statistics, including total frauds detected by each model.
- Likelihood analysis, where it is determined how many models agree on each transaction.
- Possible visualizations with probability distribution histograms and bar charts.

The analysis of the batch predictions can be summarized into the following patterns:

- The CatBoost + PCA model detected a relatively high proportion of fraudulent transactions in comparison to the other models because it is more robust with categorical features, matching the number of frauds detected to actual fraud ratio closely.
- The XGBoost + PCA model detected a few less fraudulent transactions, but the model was able to detect some fraud patterns that were missed by CatBoost, making the XGBoost model complementary to CatBoost.
- The Ensemble model looked for the best of both models, in that although it detected the most total amount of frauds and the least amount of false negatives, the discrepancies noted with the CatBoost and XGBoost models were not included within the ensemble.

Model	Fraud Probability	Prediction
CatBoost + PCA	0.3358	✓
XGBoost + PCA	0.2901	✓
Ensemble Learning	0.3130	✓

Model	Fraud Cases Detected	Total Transactions	Detection Rate
CatBoost + PCA	25,379	50,000	50.76%
XGBoost + PCA	25,423	50,000	50.85%
Ensemble Learning	25,402	50,000	50.80%

Table 3. Probability Comparison Table

Table 4. Fraud Cases Detected by Model




Figure 3. Probability Comparison Table



Figure 4. Fraud Cases Detected by Model

Batch Prediction

The Batch Prediction module facilitates the processing of multiple transactions uploaded through a CSV file. For every transaction, it computes:




Figure 5. Output images

Explanation

This screen provides the real time prediction interface for the fraud detection system. The user receives a detailed input form for each of the eight engineered transaction features. Once the user details the transaction information and submits the request, the system will run the data through the three distinct deployed machine learning models (CatBoost + PCA, XGBoost + PCA, and the Ensemble Model) and display the results immediately. The output shows both the final classification (e.g., "Legitimate" or "Fraud") and the respective Confidence Score (or fraud probability) suggested by each model, allowing for a quick, side-by-side comparison of a single transaction's risk.

This detailed screen enables models to be assessed and compared based on performance metrics from the test dataset. This dashboard provides clarity by comparing the CatBoost + PCA, XGBoost + PCA and Ensemble Learning models side-by-side. A tabular format provides a convenient way to assess key classification metrics, i.e. AUC Score, Precision, Recall and F1-Score - and bar charts provide easy directional comparison for each classification metric, so that strengths and weaknesses are easy to assess. Importantly, there is also a Radar Chart that allows for a combined multidimensional view of all performance metrics, and then a section for recommendations for which model may work better based on operational needs - it is clear that the Ensemble Model is the most robust choice for production.

The Batch Prediction screen is a convenient way to process and analyze a large number of transactions through the uploading of a .csv file. After the batch processes, this section focuses on describing the individual or combined performance of the models on the batch. The bar charts in the Prediction Summary describe the aggregate number of fraud cases detected by each model. One of the more interesting features is the Agreement and Disagreement Analysis Chart that provides a graphical representation of the amount of agreement among the three models. This chart shows how many transactions were classified as fraud or non-fraud by 0, 1, 2, or all 3 models. This is an important feature to the understanding of model certainty and for identifying the most ambiguous, possibly riskiest transactions requiring human follow-up.

CONCLUSION AND FUTURE WORK

In this work, an effective online payment fraud detection system was developed using machine learning techniques. The combination of CatBoost and XGBoost models through an ensemble approach resulted in improved predictive performance. PCA was used for dimensionality reduction, and SMOTE was applied to address class imbalance, which significantly enhanced the model's ability to detect fraudulent transactions.

The system achieved high accuracy, precision, recall, and AUC scores, demonstrating its effectiveness in real-world scenarios. Additionally, the deployment of the model using Streamlit provides a practical interface for real-time fraud detection.

Future work can focus on integrating deep learning approaches such as neural networks and graph-based models to capture more complex transaction patterns. Furthermore, real-time streaming data and largescale deployment can be explored to improve scalability and adaptability in dynamic financial environments.

REFERENCES

1. M. Habibpour, H. Gharoun, M. Mehdipour, A. Tajally, H. Asgharnezhad, A. Shamsi, A. Khosravi, M. Shafie-Khah, S. Nahavandi, and J. P. S. Catalao, "Uncertainty-aware Online payment fraud detection using deep learning 2021; arXiv:2107.13508.
2. A. Cherif, A. Badhib, H. Ammar, S. Alshehri, M. Kalkatawi, and A. Imine. "Online payment fraud detection in the era of disruptive technologies: A systematic review." J. King Saud Univ. Computer and Information Science, vol. 35, no. 1, pp. 145-174, Jan. 2023, doi:10.1016/j.jksuci.2022.11.008.
3. T. K. Dang, T. C. Tran, L. M. Tuan, and M. V. Tiep. "Machine learning based on resampling approaches and deep reinforcement learning for Online payment fraud detection systems." Appl. Sci., vol. 11, no. 21, p. 10004, Oct. 2021; doi: 10.3390/app112110004.

4. Chaquet-Ulldemolins et al., "On the black-box problem for fraud detection using machine learning (I): Linear models and informative feature selection," *Applied Sciences*, vol. 12, no. 7, p. 3328, March 2022, doi: 10.3390/app12073328.
5. E. F. Malik, K. W. Khaw, B. Belaton, W. P. Wong, and X. Chew. "Online payment fraud detection using a new hybrid machine learning architecture." *Mathematics*, vol. 10, no. 9, p. 1480, April 2022; doi: 10.3390/math10091480.
6. I. Benchaji, S. Douzi, B. El Ouahidi, and J. Jaafari, "Enhanced Online payment fraud detection using attention mechanism and LSTM deep model," *J. Big Data*, vol. 8, no. 1, p. 151, December 2021; doi: 10.1186/s40537-021-00541-8.
7. 10.1186/s40537-021-00541-8.
8. E. Esenogho, I. D. Mienye, T. G. Swart, K. Aruleba, and G. Obaido. "A neural network ensemble with feature engineering for improved Online payment fraud detection." *IEEE Access*, vol. 10, pp. 16400–16407, 2022; doi: 10.1109/ACCESS.2022.3148298.
9. E. Btoush, X. Zhou, R. Gururaian, K. Chan, and X. Tao, "A survey on Online payment fraud detection approaches in banking sector for cyber security," in *Proc. 8th International Conf. Behav. Social Comput. (BESC)*, Oct. 2021, pp. 1-7, doi: 10.1109/BESC53957.2021.9635559.
10. Y. Xie, G. Liu, C. Yan, C. Jiang, M. Zhou, and M. Li, "Learning transactional behavioral representations for Online payment fraud detection *IEEE Trans. Neural Networks and Learning Systems*, early access, October 5, 2022, doi: 10.1109/TNNLS.2022.3208967.
11. J. Yang & J. Guan "A Online payment prediction model based on feature optimization and the smote-Xgboost algorithm," *Information*, vol. 13, no. 10, p. 475, Oct. 2022, doi: 10.3390/info13100475.