

FRAUD DETECTION IN FINTECH AND DIGITAL PAYMENTS USING A HYBRID MODEL COMBINING MACHINE LEARNING AND RISK MONITORING

Do Thi Thuy Nga^{1*}, Nguyen Thi Hai Hang¹, Pham Thi Thanh Hoa¹,
Vu Van Dong², Huynh Thi Huong Thao²

¹College of Foreign Economic Relations, Ho Chi Minh City

²Ho Chi Minh City University of Industry and Trade

*Email: 6023250003@huit.edu.vn

Received: 8 April 2026; Revised: 9 May 2026; Accepted: 22 May 2026

ABSTRACT

The rapid development of Fintech and digital payments has increased the need for effective fraud detection and risk monitoring. This study evaluates fraud detection performance using imbalanced credit card transaction data. It compares Logistic Regression, Random Forest, and XGBoost, with XGBoost combined with SMOTE as the main approach. The dataset contains anonymized transaction features, where fraudulent transactions account for only a very small proportion of total observations. Model performance is assessed using Precision, Recall, F1-score, ROC-AUC, and confusion matrix analysis. The empirical results show that each model reflects a different fraud detection profile. Logistic Regression achieves high Recall but produces many false alerts. Random Forest controls false alerts more effectively and achieves the highest Precision and F1-score. XGBoost achieves the highest ROC-AUC and maintains higher Recall than Random Forest, indicating a more suitable balance for reducing missed fraud cases while controlling false alerts. The findings provide empirical evidence for the value of combining XGBoost with SMOTE in handling class imbalance. They also suggest that model outputs can be linked with Key Risk Indicators (KRIs) to support early warning and transaction risk monitoring in digital payment environments.

Keywords: Fraud detection, digital payment, XGBoost, SMOTE, risk monitoring.

1. INTRODUCTION

The development of Fintech has significantly changed banking and financial activities. According to the Financial Stability Board, Fintech refers to technology-enabled innovation in financial services. It may create new products, processes, applications, and business models [1]. In this context, e-wallets, digital banking, and online payment services have become increasingly common. These payment methods make transactions faster and more convenient, while also reducing processing costs. Digital payments have also expanded users' access to financial services [2]. However, the expansion of digital payments has also increased financial fraud risks. As the scale and speed of transactions grow, fraudulent behavior has become more sophisticated. Fraud may occur at different stages of the transaction process. Common forms include identity fraud, the use of unauthorized devices, and access from high-risk IP addresses. Fraud may also arise from the exploitation of system vulnerabilities, chargeback abuse, or unauthorized transactions. Therefore, fraud detection has become an important component of risk management for financial institutions and payment intermediaries [3].

Traditional fraud detection systems are often based on business rules. This approach is transparent, easy to implement, and easy to control. However, fixed rules are often less flexible when fraud patterns change rapidly. In digital payment environments, transaction data are usually large, fast-moving, and continuously changing. As a result, rule-based systems may generate false alerts and miss new fraudulent behavior. Machine learning is therefore considered a promising approach. It can learn from historical data and identify more complex anomalous patterns. Data mining techniques from machine learning can effectively support financial fraud detection [4].

Nevertheless, fraud detection still faces a major challenge related to class imbalance. In practice, fraudulent transactions usually account for only a very small proportion of legitimate transactions. This may cause models to become biased toward the non-fraud class. As a result, their ability to detect fraudulent transactions may decrease. To address this issue, SMOTE is used to create synthetic observations for the minority class [5]. This provides the model with more information to learn the characteristics of fraudulent transactions. XGBoost is also an effective boosting algorithm for structured data. It can capture nonlinear relationships and interactions among variables [6]. Therefore, this study combines XGBoost and SMOTE to improve fraud detection performance in digital payments.

Beyond model development, this study also considers the risk management implications of prediction results. The output of the model can be used as a transaction risk score. This score can then be combined with Key Risk Indicators (KRIs). In practice, KRIs may reflect abnormal transaction frequency, new access devices, unusual geographic locations, sudden changes in transaction value, or repeated authentication failures. Combining model-based risk scores with KRIs can support the early identification of suspicious transactions. It can also help route transactions according to risk levels, instead of relying only on a fixed alert threshold.

Based on the above discussion, this study addresses three research questions. The first question examines which model performs better in fraud detection among Logistic Regression, Random Forest, and XGBoost. The second question focuses on how the combination of XGBoost and SMOTE improves fraud detection performance. The third question considers how model results can support risk monitoring and early warning through KRIs. Accordingly, the study contributes to the literature and practice in several ways. It systematizes an approach to fraud detection in Fintech and digital payments by linking rule-based methods, machine learning, and hybrid models. It also provides empirical evidence on the effectiveness of XGBoost combined with SMOTE for imbalanced transaction data. The results are compared with Logistic Regression and Random Forest to clarify differences in model performance. Finally, the study proposes a way to connect model outputs with a risk monitoring framework through KRIs, thereby suggesting a practical direction for early warning of digital transaction risks.

2. LITERATURE REVIEW

2.1. Traditional approaches to fraud detection

In financial risk management, fraud detection is an important component of transaction risk control. In digital payments, fraud detection systems do not only identify abnormal transactions. They also support early warning, reduce losses, and protect user experience. Before machine learning was widely adopted, financial institutions mainly relied on rule-based approaches. Accordingly, traditional systems often evaluate transactions using rules derived from expert knowledge and known risk scenarios. These rules are usually formulated as “if–then” conditions, such as triggering alerts when a transaction exceeds a value threshold, occurs

in a high-risk area, or shows unusual frequency. Other signals may include repeated login failures, sudden device changes, or transactions outside normal user behavior.

The main advantages of rule-based approaches are transparency, ease of implementation, and controllability. Financial institutions can clearly explain why a transaction is flagged as risky. This is consistent with audit, compliance, and internal governance requirements. Methods such as risk scoring and alert thresholds are also used to classify transactions by risk level [3], [4]. However, traditional approaches depend heavily on fixed rules, making them less adaptive to new fraud patterns. In digital payment environments, transaction data are large, fast-moving, and continuously changing. Excessive reliance on static thresholds may increase false alerts, raise verification costs, and affect user experience [2], [4]. In addition, traditional methods face difficulties with imbalanced data, where fraudulent transactions account for only a small proportion of legitimate transactions. Rule-based systems may therefore miss complex abnormal behavior [4], [9].

2.2. Machine learning-based approaches to fraud detection

The development of big data and artificial intelligence has promoted the application of machine learning in financial fraud detection. This approach is particularly relevant to digital payment environments, where transaction data are large, fast-moving, and complex. Unlike traditional rule-based methods, machine learning models can learn from historical data and identify abnormal patterns without manually defining all rules. As a result, they can capture nonlinear relationships and interactions among multiple variables in transaction behavior. Various algorithms have been used in fraud detection, including Logistic Regression, Decision Tree, Random Forest, Support Vector Machines, Gradient Boosting, and XGBoost [4], [7]. Logistic Regression is often used as a baseline model because of its simple structure and interpretability. However, this linear model may be limited when the data contain nonlinear relationships. By contrast, Random Forest and Gradient Boosting often perform better with complex data. Random Forest uses a bagging mechanism to reduce variance and limit overfitting, while XGBoost sequentially optimizes decision trees to reduce prediction errors [6], [7].

Empirical studies have shown that machine learning often improves fraud detection compared with traditional statistical methods [7]. This advantage comes from its ability to process large datasets, detect abnormal patterns, and adapt to different transaction features. Evidence from Vietnam also suggests that Random Forest and XGBoost perform better than Logistic Regression in financial risk prediction tasks [8]. However, machine learning still faces several important challenges. Fraud data are often severely imbalanced, which may cause models to become biased toward legitimate transactions [4], [7]. In addition, many models have black-box characteristics, reducing their interpretability in financial settings. Fraud behavior may also change over time, a phenomenon known as concept drift [7]. These limitations suggest that machine learning should be combined with suitable data processing techniques and monitoring mechanisms.

2.3. Hybrid approaches to fraud detection

To address the limitations of both traditional methods and machine learning, recent studies have paid greater attention to hybrid models. In essence, a hybrid model combines several components within a single fraud detection framework. These components may include business rules, machine learning algorithms, data imbalance handling techniques, and risk monitoring mechanisms. This approach helps exploit the strengths of different methods, instead of relying on a single tool. In this framework, business rules help maintain transparency and controllability. This is important in financial environments, where alert decisions need to be clearly explained. Meanwhile, machine learning helps detect more complex fraud patterns,

especially when abnormal behavior does not follow fixed thresholds. This combination allows the system to maintain control while improving its ability to identify new risk patterns [4], [7].

Another important component of hybrid models is class imbalance handling. In fraud detection, fraudulent transactions usually account for only a very small proportion of legitimate transactions. Without proper treatment, a model may achieve high overall accuracy but still miss many fraud cases. SMOTE is used to create synthetic observations for the minority class. This improves the model's ability to learn from rare fraud cases [5]. When combined with XGBoost, this technique can help the model better identify fraud patterns in imbalanced data. Beyond classification performance, hybrid models should also be placed within a risk monitoring framework. In practice, the model output can be converted into a transaction risk score. This score can be combined with Key Risk Indicators, such as abnormal transaction frequency, new access devices, unusual geographic locations, or repeated authentication failures. As a result, the system does not only predict fraudulent transactions. It can also support early warning and case routing according to risk levels.

Although fraud detection research has made important progress, several gaps remain. Many studies still focus mainly on classification performance using transaction data. Meanwhile, the connection between models, KRIs, and risk monitoring frameworks remains limited. In addition, the combination of XGBoost, SMOTE, and risk monitoring has not been fully examined within a unified architecture. Therefore, this study adopts an XGBoost–SMOTE approach. It also discusses the potential extension of the model toward KRI integration in practical digital payment environments.

3. METHODOLOGY

3.1. Research data and data preprocessing

This study uses the Credit Card Fraud Detection dataset published on Kaggle by the Machine Learning Group, Université Libre de Bruxelles [10]. The dataset contains 284,807 credit card transactions, including 284,315 legitimate transactions and 492 fraudulent transactions. For confidentiality reasons, the dataset was anonymized before publication. It includes 28 main features transformed using PCA and labeled from V1 to V28. In addition, the dataset contains two variables, Time and Amount, which were not transformed by PCA. The target variable is Class, where Class = 0 indicates a legitimate transaction and Class = 1 indicates a fraudulent transaction.

Data preprocessing was conducted before model training. The variables Time and Amount were standardized using StandardScaler to place them on a comparable scale with the other features. The dataset was then divided into two subsets, with 80% used for training and 20% used for testing. Due to the high level of class imbalance, SMOTE was applied to the training set. SMOTE creates synthetic observations for the fraud class, helping the model learn the characteristics of the minority class more effectively. At the same time, SMOTE was applied only to the training set to avoid data leakage. The test set was kept unchanged to better reflect real classification conditions.

3.2. Model development and evaluation metrics

The empirical analysis compares three models Logistic Regression, Random Forest, and XGBoost. Among them, XGBoost combined with SMOTE is regarded as the main approach of the study. The other two models serve as benchmarks to clarify the relative effectiveness of the proposed model. Since the fraud data are highly imbalanced, Accuracy is not used as the central evaluation metric. A model may achieve high Accuracy by predicting most transactions as legitimate, but this does not properly reflect its fraud detection capability. Therefore, the study uses Precision, Recall, F1-score, ROC-AUC, and the confusion matrix. Precision

measures the proportion of truly fraudulent transactions among all transactions flagged as fraudulent. It reflects the model’s ability to control false alerts. Recall measures the proportion of actual fraudulent transactions that are correctly detected. This metric is particularly important because missed fraud cases may cause financial losses. F1-score reflects the balance between Precision and Recall. Meanwhile, ROC-AUC measures the overall ability to distinguish legitimate transactions from fraudulent ones.

In addition to these metrics, the study uses the confusion matrix to examine different types of classification errors. False Positive refers to a legitimate transaction incorrectly flagged as fraudulent. False Negative refers to a fraudulent transaction missed by the model. Analyzing these two error types helps clarify the risk management implications of each model. It also provides a basis for discussing the potential application of the model to early warning and transaction routing

4. RESULTS AND DISCUSSION

Before comparing the models, the study first conducted a correlation analysis between the explanatory variables and the target variable, Class, in order to preliminarily identify those features that exhibit notable associations with fraudulent behavior. The results are presented in Figure 1 below.




Figure 1. Correlation between the explanatory and the target variables, Class

As illustrated by the heatmap in Figure 1, several variables exhibit notable correlations with the target label, Class. In particular, V11 (0.15), V4 (0.13), and V2 (0.091) show relatively stronger positive correlations, whereas V17 (-0.33), V14 (-0.30), V12 (-0.26), and V10 (-0.22) display stronger negative correlations than the remaining features. However, most correlation coefficients remain at low to moderate levels, suggesting that no single variable is sufficient to fully explain fraudulent behavior. The variables Time and Amount show only very weak correlations with the target variable, implying that fraudulent behavior does not depend solely on transaction timing or transaction value, but is instead largely reflected in the complex interaction among multiple features. This finding further supports the use of machine learning

models, which are better able to capture nonlinear relationships and multivariate interactions than simple linear methods, as discussed in the following analysis.

Accordingly, the study proceeds to compare three machine learning models-Logistic Regression, Random Forest, and XGBoost-using Precision, Recall, F1-score, and ROC-AUC as evaluation metrics. The results reveal substantial differences in model performance, reflecting different trade-offs between fraud detection capability and the control of false alerts. Table 1 presents the overall evaluation results of the three models on the test set.

Table 1. Performance evaluation results of the models

<i>Model</i>	<i>ROC-AUC</i>	<i>Precision</i>	<i>Recall</i>	<i>F1-score</i>
Logistic Regression	0,9698	0,0581	0,9184	0,1094
Random Forest	0,9688	0,8144	0,8061	0,8103
XGBoost	0,9792	0,7311	0,8878	0,8018

From the table, it can be observed that Logistic Regression achieves the highest Recall (0.9184) among the three models, indicating that it is able to detect the majority of fraudulent transactions in the test set. However, its Precision is only 0.0581, which is substantially lower than that of the other two models, meaning that most transactions flagged as fraudulent are in fact legitimate. This suggests that although Logistic Regression is highly sensitive, such sensitivity is achieved at the cost of an extremely high false-alert rate, making the model difficult to apply directly in practice due to increased verification costs, disruption of legitimate transactions, and negative effects on customer experience. Its very low F1-score (0.1094) further confirms that the model fails to achieve the necessary balance between fraud detection capability and false-alert control.

By contrast, Random Forest demonstrates more stable and balanced performance. The model achieves the highest Precision (0.8144) among the three models, indicating that most transactions flagged as fraudulent indeed carry a high risk of fraud. At the same time, its F1-score (0.8103) is also the highest, reflecting a strong balance between alert accuracy and fraud detection capability. However, its Recall (0.8061) is lower than that of XGBoost, implying that Random Forest misses a larger number of fraudulent transactions. In the context of payment risk management, this suggests that Random Forest is more suitable when institutions prioritize minimizing false alerts and maintaining a more streamlined operational control process, while accepting a greater number of undetected anomalous transactions.

Meanwhile, XGBoost achieves the highest ROC-AUC (0.9792) among the three models and maintains a Recall of 0.8878, which is higher than that of Random Forest. These results indicate that XGBoost has the strongest overall discriminative ability between fraudulent and legitimate transactions while still preserving a high level of fraud detection performance. Although its Precision (0.7311) is lower than that of Random Forest, it remains at a reasonably strong level. This suggests that XGBoost offers a more balanced trade-off by identifying more fraudulent transactions without causing an excessive surge in false alerts. Given that, in fraud detection, the cost of missing fraudulent transactions is often greater than the cost of processing a certain number of false alerts, these findings suggest that XGBoost is a more appropriate choice for systems aiming to strengthen risk identification capacity while maintaining practical deployability.

To further clarify the differences among the models, the study proceeds to analyze the confusion matrix for each model. The results are presented as follows.




Figure 2. Confusion Matrix of the Logistic Regression Model

On the test set, the Logistic Regression model produced $TN = 55,406$, $FP = 1,458$, $FN = 8$, and $TP = 90$. Although it missed only 8 fraudulent transactions and correctly detected 90 fraud cases, it also generated a very large number of false alerts. This explains why the model achieved high Recall but very low Precision. In practice, such performance would increase verification costs, disrupt legitimate transactions, and potentially affect customer experience.




Figure 3. Confusion Matrix of the Random Forest Model

For Random Forest, the confusion matrix shows $TN = 56,846$, $FP = 18$, $FN = 19$, and $TP = 79$. These results clearly reflect the conservative nature of the model. Specifically, the number of legitimate transactions incorrectly flagged as fraudulent is extremely low, whereas the number of missed fraudulent transactions is higher. In other words, Random Forest is strongly optimized toward minimizing false alerts, thereby supporting a more stable and cost-efficient operational process. However, this conservativeness also makes the model less sensitive than XGBoost in maximizing the detection of fraud cases. This represents a typical trade-off and is also the reason why the study considers multiple evaluation metrics simultaneously.




Figure 4. Confusion Matrix of the XGBoost Model

For XGBoost, the confusion matrix records $TN = 56,832$, $FP = 32$, $FN = 11$, and $TP = 87$. Compared with Random Forest, XGBoost generates a slightly higher number of false alerts; however, in return, it correctly identifies more fraudulent transactions and reduces the number of missed fraud cases. In this sense, XGBoost still keeps false alerts at a very low level while maintaining a high fraud recovery rate. These results suggest that, although XGBoost does not achieve the highest Precision or the highest F1-score, it attains the highest ROC-AUC and demonstrates a more appropriate balance for risk management purposes in digital payments, particularly when financial institutions prioritize minimizing missed fraud cases while still controlling the volume of alerts.

Overall, from the perspective of addressing limitations in credit card fraud detection, the quantitative results show that each model handles a different aspect of the problem. Logistic Regression correctly detects 90 out of 98 fraudulent transactions, but it generates 1,458 false alerts. Random Forest reduces the number of false alerts to only 18 cases, but it misses 19 fraudulent transactions. XGBoost correctly detects 87 fraudulent transactions, misses only 11 cases, and generates 32 false alerts. These results indicate that XGBoost not only achieves the highest ROC-AUC, but also provides a better balance between reducing missed fraud cases and controlling false alerts. The combination of XGBoost and SMOTE therefore helps address the limitation caused by class imbalance, thereby improving the effectiveness of credit card fraud detection.

Another important result of the study is the analysis of feature importance in the XGBoost model, as presented below.




Figure 5. Feature Importance in the XGBoost Model

The feature importance plot shows that V14 makes the most prominent contribution, far exceeding the remaining features, followed by V4, V12, V8, V1, V13, V3, V17, V26, and V18. This result confirms that the model's decision-making process is not evenly distributed across all variables but is instead concentrated on a limited number of core features. When compared with the initial correlation analysis, variables such as V14, V12, and V17 not only exhibit notable correlations with the target variable Class but also appear among the most important features in the XGBoost model. This strengthens the reliability of the findings. However, some variables, such as V8, do not appear particularly prominent in the univariate correlation analysis yet still play an important role in the model. This suggests that XGBoost is able to exploit nonlinear relationships and interaction effects among multiple variables, going beyond what linear correlation analysis alone can reveal.

Several important observations can be drawn from the empirical results. First, fraud data in digital payment environments are highly imbalanced. Therefore, using Accuracy as the central evaluation metric is not appropriate. When fraudulent transactions account for only a very small proportion of the data, a model with high Accuracy does not necessarily perform well in detecting fraud. For this reason, this study prioritizes Precision, Recall, F1-score, and ROC-AUC. This approach is consistent with Chawla et al. [5], who emphasized the challenge of minority classes in imbalanced data. It is also consistent with Dasilas and Rigani [7], who suggested that fraud detection models should be evaluated based on their ability to identify fraudulent cases.

Second, the results show that each model reflects a different fraud detection profile. Logistic Regression achieves high Recall but very low Precision. This indicates that the model tends to broaden its alerts to recover more fraudulent transactions. However, this comes at the cost of many false alerts. Random Forest performs better in terms of Precision and F1-score, indicating stronger false-alert control. Meanwhile, XGBoost achieves the highest ROC-AUC and maintains a higher Recall than Random Forest. This suggests that XGBoost has stronger overall discriminative ability between fraudulent and legitimate transactions. This finding is consistent with Chen and Guestrin [6], who showed that XGBoost performs well with complex structured data. It is also in line with previous reviews on machine learning-based financial fraud detection [4], [7].

Third, the results suggest that XGBoost may be more suitable for early warning purposes. This model maintains a high level of fraud detection while avoiding the excessive false-alert problem observed in Logistic Regression. However, Random Forest remains a competitive model when institutions prioritize false-alert reduction. This observation is consistent with Ngai et al. [4], who argued that the choice of fraud detection techniques should be aligned with management objectives. In practice, model selection should not rely on a single metric. It should be linked to the risk appetite, operational capacity, and control objectives of each Fintech organization.

From an application perspective, the findings show that XGBoost combined with SMOTE has potential to support early warning mechanisms. Operationally, transaction data can be collected and preprocessed continuously. The model can then estimate the risk level of each transaction. Based on prediction results, financial institutions can establish different alert levels. For example, high-risk transactions may require additional verification. Medium-risk transactions may be monitored further or reviewed after processing. This approach inherits the advantages of machine learning and is consistent with risk management in digital payments [3], [7]. However, alert thresholds should not be fixed. They should be calibrated according to fraud prevalence, false-alert costs, and the risk control objectives of each institution.

In addition, the empirical scope of this study is based on anonymized transaction data. Therefore, the results mainly reflect the effectiveness of machine learning models using PCA-

transformed transaction features. The study does not directly test the role of KRIs. Nevertheless, KRIs remain important when extending the model into a digital payment risk monitoring system. In practice, KRIs can be used as additional signals for model predictions. Transaction frequency indicators may reflect a sudden increase in transaction volume within a short period. Access device indicators can help identify new devices, unusual devices, or multiple accounts using the same device. Location-based indicators can detect transactions from abnormal areas or high-risk IP addresses. In addition, sudden changes in transaction value and repeated authentication failures are also important signals. These indicators help the system identify suspicious transactions before losses occur.

When combined with machine learning models, KRIs can be used in two ways. First, they can become additional input variables if institutions have suitable operational data. Second, they can serve as a post-model monitoring layer to classify alert levels. For example, a transaction with a high risk score from XGBoost and a new access device may be assigned to an additional verification group. By contrast, a transaction with only a medium risk score may be further monitored if no abnormal KRI is observed. This combination makes the early warning system more flexible, rather than relying entirely on a fixed prediction threshold. Therefore, the role of KRIs in this study is understood as a practical extension of the model. Integrating KRIs may improve the interpretability of prediction results, support transaction routing, and enhance risk management processes. This also provides a basis for future studies to test the model on more complete operational data, including device information, location data, authentication records, user history, and transaction behavior.

5. CONCLUSION

This study evaluates fraud detection performance in digital payments by comparing Logistic Regression, Random Forest, and XGBoost on imbalanced credit card transaction data. The empirical results show that nonlinear machine learning models perform better than the linear model. Logistic Regression achieves high Recall, but it generates a very large number of false alerts. Random Forest controls false alerts more effectively and achieves the highest Precision and F1-score. Meanwhile, XGBoost achieves the highest ROC-AUC and maintains a higher Recall than Random Forest. These results suggest that XGBoost is more suitable when the objective is to reduce missed fraud cases while still controlling false alerts.

From a theoretical perspective, this study adds evidence to the hybrid approach to fraud detection. This approach combines machine learning, class imbalance handling, and a risk management perspective. The findings show that fraud detection models should not be evaluated only by Accuracy. Instead, Precision, Recall, F1-score, and ROC-AUC better reflect the ability to identify fraudulent transactions. The study also highlights the need to consider both predictive performance and the risk management implications of the model.

From a practical perspective, the findings suggest that XGBoost combined with SMOTE can serve as a foundation for early warning systems in digital payments. The model output can be converted into a transaction risk score, which supports case routing according to risk levels. High-risk transactions may require additional verification. Medium-risk transactions may be further monitored before a final decision is made. When combined with KRIs such as abnormal transaction frequency, new access devices, unusual geographic locations, sudden changes in transaction value, or repeated authentication failures, the system can generate more flexible alerts. This approach reduces dependence on a fixed alert threshold.

However, the study still has several limitations. The dataset used in this study was anonymized, so the original transaction variables cannot be directly interpreted. Therefore, the study has not fully tested the role of KRIs in a real operational environment. In addition, the analysis focuses on credit card transaction data and has not fully examined changes in fraud

behavior over time. Future studies may extend this work by integrating operational data on devices, locations, authentication, and user history. They may also combine model explainability methods to improve prediction transparency. Finally, future research should test the model in real-time monitoring settings to evaluate its adaptability to new fraud patterns.

REFERENCES

- [1] Financial Stability Board, “FinTech and market structure in financial services: Market developments and potential financial stability implications,” Feb. 2019. [Online]. Available: <https://www.fsb.org/uploads/P140219.pdf>
- [2] World Bank Group, “Payment systems worldwide: A snapshot,” 2020. [Online]. Available: <https://documents1.worldbank.org/curated/en/115211594375402373/pdf/A-Snapshot.pdf>
- [3] Basel Committee on Banking Supervision, Bank for International Settlements, “Digital fraud and banking: Supervisory and financial stability implications,” Nov. 2023. [Online]. Available: <https://www.bis.org/bcbs/publ/d558.pdf>
- [4] European Central Bank and European Banking Authority, “2024 report on payment fraud,” Aug. 2024. [Online]. Available: <https://www.ecb.europa.eu/press/intro/publications/pdf/ecb.ebaecb202408.en.pdf>
- [5] E. W. T. Ngai, Y. Hu, Y. H. Wong, Y. Chen, and X. Sun, “The application of data mining techniques in financial fraud detection: A classification framework and an academic review of literature,” *Decision Support Systems*, vol. 50, no. 3, pp. 559–569, Feb. 2011, doi: <https://doi.org/10.1016/j.dss.2010.08.006>
- [6] L. H. Aros, L. X. B. Molano, F. Gutierrez-Portela, J. J. M. Hernandez, and M. S. R. Barrero, “Financial fraud detection through the application of machine learning techniques: A literature review,” *Humanities and Social Sciences Communications*, vol. 11, no. 1, Art. no. 1130, 2024, doi: <https://doi.org/10.1057/s41599-024-03606-0>
- [7] T. Chen and C. Guestrin, “XGBoost: A scalable tree boosting system,” in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, San Francisco, CA, USA, 2016, pp. 785–794, doi: <https://doi.org/10.1145/2939672.2939785>
- [8] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, “SMOTE: Synthetic minority over-sampling technique,” *Journal of Artificial Intelligence Research*, vol. 16, pp. 321–357, 2002, doi: <https://doi.org/10.1613/jair.953>
- [9] T. T. D. Truong and H. T. Le, “Machine learning based bankruptcy prediction of Vietnam companies,” *Journal of Economics & Development*, no. 310, pp. 44–53, Apr. 2023, doi: <https://doi.org/10.33301/JED.VI.1066>
- [10] Machine Learning Group, Université Libre de Bruxelles, “Credit card fraud detection,” *Kaggle*, 2016. [Online]. Available: <https://www.kaggle.com/datasets/mlg-ulb/creditcardfraud>