SvedbergOpen
DISSEMINATION OF KNOWLEDGE

International Journal of Artificial Intelligence and Machine Learning

Publisher's Home Page: https://www.svedbergopen.com/

Research Paper Open Access

Adaptive Machine Learning for Detecting Emerging Financial Fraud Patterns in Digital Payment Networks

Parth Ghaswala

Cloud Infrastructure Engineer Gainwell Technologies Rancho Cordova, CA, USA. E-mail: parth4excel@gmail.com

Abstract
This study aimed to evaluate static and adaptive machine-learning models for detecting fraud in digital payment transactions under conditions of severe class imbalance and evolving fraud behaviour. The study compared Random Forest, conventional XGBoost, and Adaptive XGBoost across multiple performance measures, and further assessed static and adaptive XGBoost across five sequential evaluation periods to examine temporal robustness under a rolling-window learning strategy applied to a large simulated transaction dataset. Random Forest and XGBoost achieved high overall accuracy but limited recall, reflecting weak minority-class detection despite strong aggregate performance. Adaptive XGBoost recorded lower accuracy and precision but achieved higher recall and stronger discriminatory performance. Sequential evaluation showed static XGBoost maintained more stable precision, while Adaptive XGBoost achieved higher recall in most periods, with fluctuating precision but a stronger overall balance between precision and recall. Adaptive XGBoost demonstrated greater sensitivity to emerging fraud patterns and improved recall over time, whereas static XGBoost offered more consistent precision, indicating a trade-off that should guide model selection in real-world fraud-monitoring systems requiring both detection sensitivity and operational reliability.

Keywords: Adaptive machine learning, Class imbalance, Digital payments, Fraud detection, Gradient boosting

1. Introduction

Digital payment systems have expanded rapidly over the past decade, driven by the growth of e-commerce, mobile banking, and contactless transaction technologies. This expansion has been accompanied by a parallel rise in financial fraud, as criminals continuously adapt their methods to exploit vulnerabilities in payment infrastructure (Al-dahasi *et al*., 2025). Traditional rule-based fraud-detection systems, which rely on fixed thresholds and predefined logic, have increasingly struggled to keep pace with the speed and sophistication of modern fraud schemes (Vanini *et al*., 2023). As a result, machine learning has become the dominant approach for identifying fraudulent transactions, offering the ability to learn complex, non-linear relationships between transaction features and fraudulent behaviour rather than depending on static, manually defined rules. This shift has been documented across multiple payment contexts, including mobile banking and fintech platforms (Ononiwu *et al*., 2023), virtual payment systems (Sethupathy, 2024), and digital banking channels more broadly (Khatib, 2024).

Despite these advances, fraud detection remains a uniquely challenging classification problem. Fraudulent transactions typically represent a very small fraction of total transaction volume, creating a severe class-imbalance problem that can bias models toward the majority legitimate class and inflate accuracy metrics while masking poor detection of the minority fraud class (Oliveira, 2026). Several studies have proposed imbalance-mitigation techniques to address this issue, including automatic algorithmic tuning approaches (Kolodiziev *et al*., 2020) and online payment fraud models that explicitly account for skewed class distributions (Almazroi and Ayub, 2023). Beyond imbalance, fraud behaviour is inherently non-stationary: fraudulent actors continuously modify their tactics to evade detection, meaning that patterns learned from historical data can quickly become outdated. Adversarial behaviour further compounds this problem, as fraudsters may deliberately manipulate transaction features to evade high-frequency detection systems (Alam *et al*., 2025).

To address this limitation, researchers have increasingly explored adaptive machine-learning frameworks capable of updating themselves in response to new data. Integrated drift-detection and adaptive-windowing techniques have been proposed to help fraud-detection pipelines respond to evolving transaction behaviour in near real time (Al Lawati *et al*., 2025). Adaptive strategies have also been extended to payment-gateway security, where resilient, self-updating models are used to counter evolving cyber threats (Karangara, 2025), and to edge-based financial risk assessment, where hybrid deep learning and adaptive regression models

support low-latency decision-making (Ahmed *et al*., 2025). Deep neural network architectures have likewise been applied within adaptive learning frameworks to strengthen fraud-detection robustness (Enyiorji *et al*., 2025), while broader infrastructure-level frameworks have combined machine learning and AI-based adaptive detection to support resilient digital payment ecosystems (Young, 2025).

Related work has also examined complementary strategies for improving fraud-detection robustness, including decentralised machine-learning architectures designed to balance detection effectiveness with data-privacy compliance (AbouGrad and Sankuru, 2025). Anomaly-detection and real-time monitoring approaches have further been proposed to secure digital transactions as they occur (Mistry *et al*., 2024), while systematic reviews of e-commerce fraud-detection literature confirm that machine-learning techniques now dominate the field, though with persistent challenges around generalizability and imbalance (Mutemi and Bacao, 2024).

Ensemble tree-based methods, particularly Random Forest and XGBoost, have become especially prominent in fraud-detection research due to their strong performance on structured tabular data and their ability to capture complex feature interactions without extensive manual feature engineering. Adaptive extensions of these ensemble methods have recently been evaluated specifically for transaction-level fraud detection in digital payment settings (Mahalakshmi *et al*., 2026) and within FinTech environments more broadly,while adaptive system-level frameworks have also been proposed to help institutions identify and respond to continuously evolving financial threats (Ikemefuna *et al*., 2024). However, most existing evaluations of these models report performance on a single, static train-test split, which does not reveal how detection performance changes as transaction behaviour evolves over successive time periods(Upadhyay, 2026).

The present study addresses this gap by evaluating and comparing the performance of Random Forest, conventional XGBoost, and an adaptive version of XGBoost using a simulated digital payment transaction dataset spanning a 183-day observation period with over 1.7 million transactions, incorporating realistic class imbalance and time-dependent fraud scenarios.

The objective of this study is to evaluate and compare the performance of Random Forest, conventional XGBoost, and Adaptive XGBoost in detecting fraudulent transactions within a highly imbalanced digital payment dataset, using multiple complementary performance measures including accuracy, precision, recall, F1-score, ROC-AUC, and PR-AUC, to determine how effectively each model identifies the minority fraud class beyond what aggregate accuracy alone can reveal. The temporal robustness of static and adaptive XGBoost models by examining their performance across five sequential evaluation periods, thereby determining whether an adaptive, rolling-window learning strategy can maintain or improve fraud-detection effectiveness as transaction patterns change over time, compared to a conventional static model trained on a single historical dataset.

2. Materials and Methods

2.1 Data Source

The study used the Fraud Detection Handbook Simulated Transaction Dataset to examine machine-learning approaches for identifying fraudulent digital payment activities. The dataset represents transactions generated over 183 days and contains 1,754,155 individual records. It includes 5,000 customers, 10,000 payment terminals, three defined fraud scenarios, and nine transaction variables. The dataset contains 1,739,474 legitimate transactions and 14,681 fraudulent transactions, resulting in a fraud proportion of 0.837%. Transaction amounts showed a mean of 53.63 and a median of 44.64, while the maximum recorded amount was 2,628.00. The dataset was selected because its structure incorporates both class imbalance and time-dependent fraud behaviour.

2.2 Data Characteristics

The transaction data were examined according to their legitimate and fraudulent classes before model evaluation. The marked difference between the number of legitimate and fraudulent observations demonstrates a strongly imbalanced classification setting. Fraudulent records represented only 0.837% of all transactions, whereas legitimate transactions constituted the overwhelming majority. Such a distribution was retained as part of the experimental setting because the dataset was specifically designed to represent realistic challenges associated with fraud identification. The analysis also considered the temporal organisation of the transactions, allowing model performance to be examined across successive periods rather than relying exclusively on an overall evaluation. This structure supported assessment of model behaviour under changing fraud patterns.

2.3 Machine-Learning Models

Three machine-learning approaches were considered for comparative fraud-detection analysis: Random Forest, conventional XGBoost, and Adaptive XGBoost. Random Forest and XGBoost were evaluated as reference models, while Adaptive XGBoost was examined for its ability to respond to changing transaction patterns. Model performance was assessed using several classification measures rather than accuracy alone. These measures included precision, recall, F1-score, ROC-AUC, and PR-AUC. This multi-metric evaluation was particularly relevant to the dataset because fraudulent transactions formed a relatively small proportion of the complete dataset. The comparative analysis therefore focused on how effectively each approach identified fraudulent observations while maintaining discrimination between legitimate and fraudulent transactions.

2.4 Adaptive Learning Strategy

The adaptive analysis was designed around the time-dependent nature of the transaction data. Instead of treating all observations as a single undifferentiated dataset, the study evaluated model behaviour over consecutive chronological periods. Adaptive XGBoost was assessed using a rolling-window strategy, allowing its performance to be examined as transaction patterns changed across the observation period. This approach enabled comparison with static XGBoost under sequential evaluation conditions. The method was particularly relevant to the study objective because emerging fraud behaviour may vary between different periods. The sequential design therefore provided a basis for examining whether an adaptive model could maintain fraud-detection performance when evaluated on transactions occurring at different points in time.

2.5 Sequential Evaluation Periods

For temporal assessment, the dataset was divided into five consecutive evaluation periods: August 29–September 4, September 5–September 11, September 12–September 18, September 19–September 25, and September 26–September 30. The corresponding transaction volumes ranged from 47,759 to 67,313, while the number of fraud cases ranged from 412 to 627. Static and Adaptive XGBoost were compared within each period using precision, recall, F1-score, and PR-AUC. Mean values were subsequently reported across the evaluation periods. This chronological procedure enabled the study to examine fluctuations in model performance and provided a direct comparison between conventional static learning and adaptive learning under sequential transaction conditions.

2.6 Performance Evaluation

Model effectiveness was assessed using complementary performance indicators. Accuracy was reported for the overall comparison of Random Forest, XGBoost, and Adaptive XGBoost, while precision, recall, and F1-score were used to examine fraud-class identification performance. ROC-AUC was included to evaluate overall discriminatory capability, whereas PR-AUC provided an additional measure appropriate for the imbalanced transaction setting. For the sequential analysis, precision, recall, F1-score, and PR-AUC were calculated for both static and adaptive XGBoost across each evaluation period. Mean performance values were also obtained to provide an overall temporal comparison. The reported results were derived from the authors' analysis of the Fraud Detection Handbook simulated transaction dataset.

3. Results

3.1 Characteristics of the Simulated Digital Payment Transaction Dataset

This table presents the overall characteristics of the simulated digital payment transaction dataset used for the analysis. The dataset contains 1,754,155 transactions recorded over an observation period of 183 days, providing a substantial basis for evaluating fraud-detection models. Among these transactions, 1,739,474 were identified as legitimate, whereas 14,681 were fraudulent, corresponding to a fraud rate of 0.837% as shown in Table 1. The data represent transactions involving 5,000 customers and 10,000 payment terminals across three fraud scenarios. Nine transaction variables were included in the dataset. The mean transaction amount was 53.63, while the median was 44.64, with the maximum transaction amount reaching 2,628.00. These characteristics indicate a highly imbalanced transaction environment.

Table 1. Characteristics of the Simulated Digital Payment Transaction Dataset

Characteristic	Result
Total transactions	1,754,155
Legitimate transactions	1,739,474
Fraudulent transactions	14,681

Fraudulent transactions (%)	0.837
Observation period	183 days
Customers	5,000
Payment terminals	10,000
Fraud scenarios	3
Transaction variables	9
Mean transaction amount	53.63
Median transaction amount	44.64
Maximum transaction amount	2,628.00

Source: Le Borgne *et al*. (2022)

3.2 Comparative Performance of Machine-Learning Models for Fraud Detection

This table presents the fraud-detection performance of Random Forest, XGBoost, and Adaptive XGBoost using multiple evaluation measures. Random Forest and conventional XGBoost achieved the same reported accuracy of 99.293%, while their precision, recall, F1-score, ROC-AUC, and PR-AUC values differed, as shown in Table 2. Random Forest recorded 100.000% precision, whereas XGBoost achieved 77.528%. Both models reported a recall of 29.320% and an F1-score of 42.549%. Adaptive XGBoost produced a lower accuracy of 98.935% but achieved the highest recall at 30.559% and the highest ROC-AUC at 0.729. Its PR-AUC of 0.309 was also slightly higher than the corresponding value for the other models, suggesting improved discrimination under the imbalanced fraud-detection setting.

Table 2. Comparative Performance of Machine-Learning Models for Fraud Detection

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-score (%)	ROC-AUC	PR-AUC
Random Forest	99.293	100.000	29.320	42.549	0.587	0.305
XGBoost	99.293	77.528	29.320	42.549	0.655	0.305
Adaptive XGBoost	98.935	38.051	30.559	33.896	0.729	0.309

Source: Le Borgne *et al*. (2022)

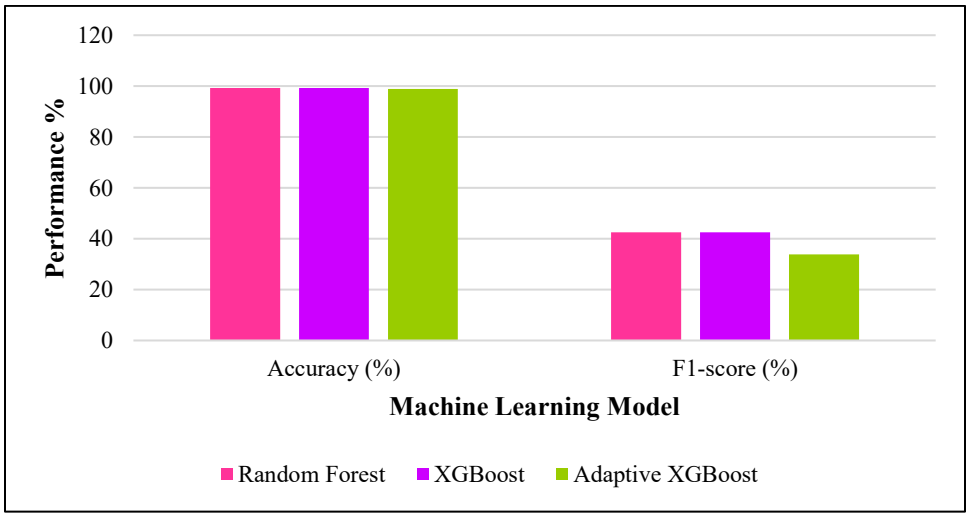

Figure 2. Comparative Performance of Machine Learning Models

The figure shows the comparative performance of Random Forest, XGBoost, and Adaptive XGBoost across the principal classification measures. Random Forest and XGBoost achieved identical accuracy values of 99.293%, while Adaptive XGBoost recorded a slightly lower accuracy of 98.935%, as shown in Figure 2. Random Forest produced the highest precision at 100.000%, followed by XGBoost at 77.528% and Adaptive XGBoost at 38.051%. In contrast, Adaptive XGBoost achieved the highest recall at 30.559%, exceeding the 29.320% recorded by both conventional models. The F1-score values were 42.549% for Random Forest and XGBoost and 33.896% for Adaptive XGBoost. These findings demonstrate differences between overall accuracy and fraud-class detection performance.

3.3 Performance of Static and Adaptive XGBoost Across Sequential Evaluation Periods

This section evaluates static and adaptive XGBoost models across successive chronological periods, highlighting changes in fraud-detection performance over time. Across the five evaluation periods, the static model generally maintained higher precision, while the adaptive approach achieved higher recall during most periods, as shown in Table 3. The adaptive model's precision varied considerably, ranging from 23.02% to 69.41%, whereas its recall ranged from 28.26% to 34.22%. The highest adaptive PR-AUC was 0.359 during September 26–30, compared with 0.345 for the static model during the same period. Overall, the mean results show 50.66% precision, 30.54% recall, 37.50% F1-score, and 0.315 PR-AUC for adaptive XGBoost, compared with 79.30%, 29.55%, 42.78%, and 0.309 for static XGBoost.

Table 3. Performance of Static and Adaptive XGBoost Across Sequential Evaluation Periods

Evaluation Period	Transactions	Fraud Cases	Static Precision (%)	Adaptive Precision (%)	Static Recall (%)	Adaptive Recall (%)	Static F1-score (%)	Adaptive F1-score (%)	Static PR-AUC	Adaptive PR-AUC
Aug 29–Sep 4	66,670	627	84.78	53.48	28.07	28.87	40.46	38.84	0.291	0.304
Sep 5–Sep 11	67,259	580	73.42	69.41	32.07	30.86	44.71	42.87	0.329	0.323
Sep 12–Sep 18	67,196	607	81.34	43.78	29.32	30.48	43.15	35.99	0.307	0.302
Sep 19–Sep 25	67,313	598	79.71	23.02	25.75	28.26	39.04	25.30	0.272	0.288
Sep 26–Sep 30	47,759	412	77.27	63.60	32.52	34.22	46.53	44.48	0.345	0.359
Mean	63,239	565	79.30	50.66	29.55	30.54	42.78	37.50	0.309	0.315

Source: Le Borgne *et al*. (2022)

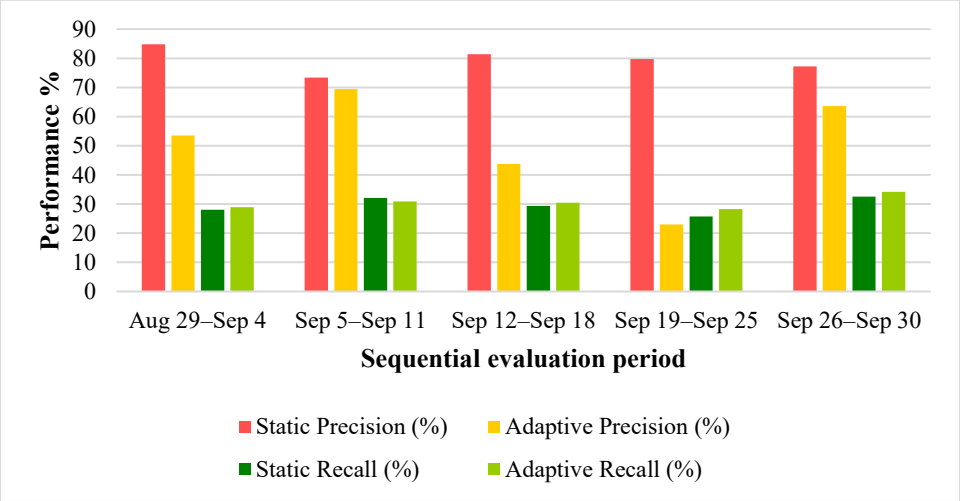

Figure 3. Static and Adaptive XGBoost Performance Across Evaluation Periods

This figure compares the static and adaptive XGBoost models' performance across five consecutive evaluation periods. Static XGBoost consistently produced higher precision than the adaptive model, with precision values ranging from 73.42% to 84.78% as shown in Figure 3. Adaptive XGBoost showed greater variation, reaching its highest precision of 69.41% during September 5–11 and its lowest value of 23.02% during September 19–25. The adaptive model recorded higher recall than the static approach in most periods, with its highest recall of 34.22% occurring during September 26–30. The sequential results indicate that model performance changed across evaluation periods, reflecting the temporal variability represented within the simulated transaction data.

4. Discussion

The study provides meaningful insight into how conventional and adaptive machine-learning approaches perform when applied to a highly imbalanced digital payment dataset. Although Random Forest and XGBoost achieved near-identical accuracy of 99.293%, this figure alone does not reflect their true fraud-detection capability, given that fraudulent transactions represented only 0.837% of the dataset. In such a skewed setting, high accuracy is easily achieved simply by correctly classifying the overwhelming majority of legitimate transactions, while genuinely fraudulent cases may still be missed at a high rate. This is confirmed by the recall values recorded for both models (29.320%), meaning that roughly seven out of ten fraudulent transactions went undetected despite the seemingly strong overall accuracy. This pattern echoes long-standing concerns in the fraud-detection literature that accuracy is a misleading metric under class imbalance, and that precision, recall, and PR-AUC provide a more honest picture of model usefulness (Chang *et al*., 2024).

The comparison between static and adaptive XGBoost adds further nuance to this interpretation. Adaptive XGBoost recorded lower overall accuracy (98.935%) and a considerably lower precision (38.051%) than the static models, yet it achieved the highest recall (30.559%), the highest ROC-AUC (0.729), and a marginally higher PR-AUC (0.309) in the overall comparison. This trade-off suggests that the adaptive model was more willing to flag borderline or unusual transactions as suspicious, which increased false positives and lowered precision, but also allowed it to capture a larger share of true fraud cases. This behaviour is consistent with the underlying logic of adaptive learning: because the model is periodically updated using recent transaction windows, it remains more responsive to shifting fraud tactics rather than relying on patterns learned from a single, static training period. This directly supports the argument made by Odunaike (2025), who emphasised that adaptive, time-series-based fraud models are better suited to dynamic risk environments than static frameworks that assume fraud behaviour remains constant over time.

The sequential evaluation results reinforce this interpretation. Static XGBoost maintained relatively stable and generally higher precision across all five periods (ranging from 73.42% to 84.78%), whereas Adaptive XGBoost's precision fluctuated sharply, dropping to as low as 23.02% during September 19–25. At the same time, Adaptive XGBoost generally produced higher recall across the periods, and its PR-AUC exceeded that of the static model in the final period (0.359 versus 0.345). This fluctuation illustrates a key vulnerability of adaptive models: because they continually retrain on recent data, they are sensitive to short-term noise and shifts in the rolling window, which can temporarily degrade precision even as recall improves. This volatility aligns with observations made by Pranto *et al*. (2022), who noted that adaptive, incentive-based fraud-detection systems can experience performance instability when transaction patterns shift rapidly, even while offering better long-term responsiveness than fixed models.

From a practical standpoint, these findings carry direct implications for real-world fraud-prevention systems. Khurana (2020) argued that real-time transaction security depends not only on detection accuracy but on a model's ability to adapt to emerging fraud typologies without requiring constant manual retraining; the present findings support this view, since the adaptive model's recall advantage suggests better sensitivity to newly emerging fraud behaviour. However, the precision instability observed here also aligns with the caution raised by Chang *et al*. (2022), who noted that overly reactive fraud-detection systems risk generating excessive false alerts, which can burden compliance teams and erode customer trust if not carefully calibrated. Similarly, Al-dahasi et al. (2025) demonstrated that imbalance-mitigation strategies must be paired with adaptive mechanisms to avoid sacrificing precision for recall gains, a balance that this study's adaptive model has not yet fully achieved.

This study has several limitations. First, the analysis relied on a simulated dataset, which, despite incorporating realistic imbalance and temporal fraud structures, may not fully capture the complexity of real-world payment behaviour, customer heterogeneity, or adversarial evasion tactics. Second, the adaptive model's rolling-window configuration was not systematically varied, so its sensitivity to window size and update frequency remains unexplored. Third, the evaluation period spanned only 183 days, which may be too short to observe longer-term seasonal or cyclical fraud patterns. Finally, the study focused on tree-based ensemble methods only, without benchmarking against deep learning or hybrid adaptive-ensemble approaches.

Future research should test the adaptive framework on real, anonymised transaction data to validate its generalizability. Further work could also explore adaptive threshold-tuning or cost-sensitive learning to stabilise precision without sacrificing the recall gains observed here, along with hybrid architectures combining adaptive XGBoost with deep learning or drift-detection mechanisms to better balance responsiveness and reliability in production fraud-monitoring systems.

5. Conclusion

This study compared static and adaptive machine-learning approaches for detecting fraud within a large, highly imbalanced digital payment dataset containing over 1.7 million transactions, of which only 0.837% were fraudulent. Random Forest and conventional XGBoost achieved high overall accuracy but showed limited recall, meaning a large share of fraudulent transactions remained undetected despite strong aggregate performance. Adaptive XGBoost, evaluated through a rolling-window strategy across five sequential periods, produced lower precision and accuracy but consistently higher recall and PR-AUC, indicating a greater capacity to identify fraudulent transactions as patterns changed over time. The sequential analysis further showed that the adaptive model's performance fluctuated more than the static model's, particularly in precision, reflecting its sensitivity to short-term shifts in transaction behaviour. Taken together, these findings indicate that no single model configuration is optimal across all evaluation criteria. Static models offer more stable precision, while adaptive models are better positioned to respond to emerging fraud typologies, at the cost of increased false positives. This trade-off has direct relevance for financial institutions designing fraud-monitoring systems, as it highlights the need to balance detection sensitivity against operational costs associated with false alerts. Future work should validate these results using real transaction data and explore mechanisms that stabilise adaptive precision while preserving its demonstrated advantage in recall and temporal responsiveness.

References

1. AbouGrad, H., & Sankuru, L. (2025). Online banking fraud detection model: Decentralised machine learning framework to enhance effectiveness and compliance with data privacy regulations. Mathematics, 13(13), 2110.
2. Ahmed, M. P., Tisha, S. A., & Sweet, M. R. (2025). Real-Time Hybrid Optimisation Models for Edge-Based Financial Risk Assessment: Integrating Deep Learning with Adaptive Regression for Low-Latency Decision Making. Journal of Business and Management Studies, 7(7), 38-52.
3. Al Lawati, H. M., Zainal, A., Al-Rimy, B. A. S., Al-Azawi, M., Kassim, M. N., Almalki, S. A., & Alghamdi, T. A. (2025). An integrated preprocessing and drift detection approach with adaptive windowing for fraud detection in payment systems. IEEE Access, 13, 92036-92056.
4. Alam, M. K., Mahmud, M. A., & ALAM, M. A. (2025). Adversarial machine learning for robust fraud detection in high-frequency financial transactions. Journal of Computer Science and Technology Studies, 7(8), 314-335.
5. Al-dahasi, E. M., Alsheikh, R. K., Khan, F. A., & Jeon, G. (2025). Optimizing fraud detection in financial transactions with machine learning and imbalance mitigation. Expert Systems, 42(2), e13682.
6. Almazroi, A. A., & Ayub, N. (2023). Online payment fraud detection model using machine learning techniques. IEEE Access, 11, 137188-137203.
7. Chang, V., Ali, B., Golightly, L., Ganatra, M. A., & Mohamed, M. (2024). Investigating credit card payment fraud with detection methods using advanced machine learning. Information, 15(8), 478.
8. Chang, V., Di Stefano, A., Sun, Z., & Fortino, G. (2022). Digital payment fraud detection methods in digital ages and Industry 4.0. Computers and Electrical Engineering, 100, 107734.
9. Enyiorji, P., Oloke, K., & Ogundipe, A. (2025). Adaptive learning architectures for financial fraud detection: A deep neural network approach. Int. J. Financ. Manag. Econ, 8(2), 1031-1042.
10. Ikemefuna, C. D., Okusi, O., Iwuh, A. C., & Yusuf, S. (2024). Adaptive fraud detection systems: Using ML to identify and respond to evolving financial threats. International Research Journal of Modernisation in Engineering, 6, 2077-2092.
11. H. K. Mistry, A. M. Goswami, and C. C. Mavani, "Automated anomaly detection and response system for enhancing cloud security," Indian Patent Application, Jul. 2024.
12. Karangara, R. (2025). Adaptive machine learning models for securing payment gateways: A resilient approach to mitigating evolving cyber threats in digital transactions. Artificial Intelligence Evolution, 44-64.
13. Khatib, S. (2024). The application of machine learning models in fraud detection and prevention across digital banking channels and payment platforms. International Journal of Advanced Computational Methodologies and Emerging Technologies, 14(9), 1-7.
14. Khurana, R. (2020). Fraud detection in ecommerce payment systems: The role of predictive AI in real-time transaction security and risk management. International Journal of Applied Machine Learning and Computational Intelligence, 10(6), 1-32.
15. Kolodiziev, O., Mints, A., Sidelov, P., Pleskun, I., & Lozynska, O. (2020). Automatic machine learning algorithms for fraud detection in digital payment systems. Восточно-Европейский журнал передовых технологий, 5(9-107), 14-26.

16. Le Borgne, Y. A., Siblini, W., Lebichot, B., & Bontempi, G. (2022). Reproducible Machine Learning for Credit Card Fraud detection: Practical handbook.
17. Mahalakshmi, S., Hemamalini, R. R., Alfred, M. C., & Ramya, K. (2026, March). Transaction Fraud Detection in Digital Payments Using Adaptive Learning Models. In 2026 International Conference on Data Science, Agents and Artificial Intelligence (ICDSAAI) (Vol. 1, pp. 1-6). IEEE.
18. Mutemi, A., & Bacao, F. (2024). E-commerce fraud detection based on machine learning techniques: Systematic literature review. Big Data Mining and Analytics, 7(2), 419-444.
19. Odunaike, A. (2025). Designing adaptive compliance frameworks using time series fraud detection models for dynamic regulatory and risk management environments. Comput. Eng. Appl, 15, 49-68.
20. Oliveira, D. N. O. (2026). Neural networks for real-time financial fraud detection. Journal International Review of Research Studies, 1(03), 1-49.
21. Ononiwu, M., Azonuche, T. I., Okoh, O. F., & Enyejo, J. O. (2023). Machine learning approaches for fraud detection and risk assessment in mobile banking applications and fintech solutions. International Journal of Scientific Research in Science, Engineering and Technology, 10(4), 1-10.
22. Pranto, T. H., Hasib, K. T. A. M., Rahman, T., Haque, A. B., Islam, A. N., & Rahman, R. M. (2022). Blockchain and machine learning for fraud detection: A privacy-preserving and adaptive incentive-based approach. Ieee Access, 10, 87115-87134.
23. Sethupathy, U. K. A. (2024). Fraud Detection Mechanisms in Virtual Payment Systems. International Journal of Computer Technology and Electronics Communication, 7(2), 8504-8515.
24. Upadhyay, N. (2026). Adaptive Machine Learning for Emerging Fraud Patterns in FinTech Transactions. Journal of Artificial Intelligence and Semiconductor Integrated Hardware, 1(1), 7-13.
25. Vanini, P., Rossi, S., Zvizdic, E., & Domenig, T. (2023). Online payment fraud: from anomaly detection to risk management. Financial Innovation, 9(1), 66.
26. Young, J. (2025). Building Resilient Digital Payment Infrastructure with Machine Learning and AI-Based Adaptive Fraud Detection Frameworks. Available at SSRN 5359364.