FUDMA Journal of Sciences (FJS)
ISSN online: 2616-1370
ISSN print: 2645 - 2944
Vol. 10 No. 13, July, 2026, pp 244 – 250
DOI: https://doi.org/10.33003/fjs-2026-1013-5397

An Adaptive Temporal Convolutional Network Framework for Real-Time Credit Card Fraud Detection in Highly Imbalanced Financial Transaction Data

***Richard Daniel, Etemi Joshua Garba, Yusuf Musa Malgwi and Ibrahim Ishaku**

Department of Computer Science, Modibbo Adama University, Yola, Adamawa State, Nigeria.

*Corresponding authors' email: bydanrich@gmail.com

ABSTRACT

This paper focuses on credit card fraud detection using Temporal Convolutional Networks (TCN), real-time fraud detection, imbalanced data handling, and comparative performance against Logistic Regression, Random Forest, and XGBoost. The rapid expansion of digital payments has increased the need for fraud detection systems that can operate with low latency, maintain high predictive performance, and provide interpretable decisions for financial analysts. This study develops an adaptive TCN-based framework for real-time credit card fraud detection using a highly imbalanced transaction dataset. The framework integrates preprocessing, feature engineering, chronological validation, imbalance-aware learning, comparative baseline modelling, low-latency inference testing, and SHAP-based explainability. The experiment used the Kaggle European credit card fraud dataset containing 284,807 transactions, of which 492 were fraudulent, representing a fraud rate of 0.173 percent and an imbalance ratio of 577.88:1. Results show that Random Forest achieved the highest PR-AUC of 0.773 and precision of 0.925, XGBoost achieved the highest ROC-AUC of 0.977 with p99 latency of 6.40 ms, while the TCN achieved competitive fraud recall of 0.808, PR-AUC of 0.714, ROC-AUC of 0.974, and p99 latency of 6.34 ms. SHAP explanations identified stable fraud-predictive patterns across models, supporting transparent fraud investigation. The findings demonstrate that adaptive, explainable, and low-latency fraud detection frameworks can strengthen financial transaction monitoring in modern digital payment environments.

Received: 29 June 2026

Accepted: 17 July 2026

Published: 18 Aug. 2026

Keywords: Credit Card Fraud Detection; Temporal Convolutional Network; Real-Time Fraud Detection; Imbalanced Data; SHAP; Xgboost; Random Forest

INTRODUCTION

Digital payment systems have become central to modern financial activity because they support electronic transfer of monetary value through credit cards, mobile wallets, online banking, and point-of-sale platforms. The shift from cash-based systems to digital payment ecosystems has improved transaction convenience, commercial speed, and financial inclusion. Teker et al. (2022) described digital payment systems as electronic platforms that facilitate financial exchange without physical cash, while Ibrahim et al. (2024) linked the rise of cashless transactions to consumer preference for convenience and accessibility. Lakhotia (2024) similarly argued that digital payment systems are transforming global commerce by reducing dependence on traditional banking procedures.

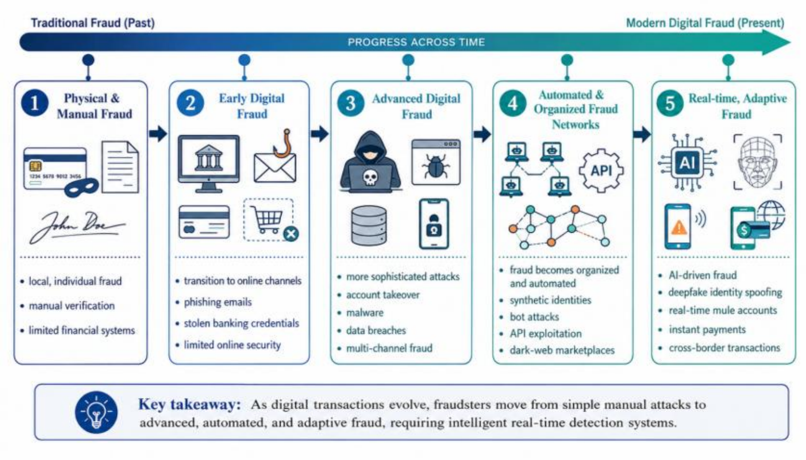

Figure 1: Evolution of Financial Fraud

The growth of digital payment channels has also expanded opportunities for fraud actors. Credit card fraud now involves unauthorised use of cardholder information through phishing, data breaches, identity spoofing, malware attacks, synthetic identities, and automated transaction manipulation. Mohammad and Logeshwaran (2024) observed that credit card fraud remains persistent because payment cards are widely used for online and physical transactions. Pasupuleti (2025) and Taralkar (2025) further emphasized that fraudulent transaction anomalies can emerge within milliseconds, making manual review unsuitable for real-time payment environments.

Traditional fraud detection systems have relied on rule-based checks and conventional machine learning models. Rule-based systems apply predefined thresholds such as unusual transaction amount, location inconsistency, or purchase frequency. Although these methods are interpretable, they are limited when fraud patterns evolve beyond fixed rules. Conventional classifiers also struggle with highly imbalanced transaction data, where legitimate transactions vastly outnumber fraudulent ones. Al-Dahasi et al. (2024) noted that such imbalance can bias models toward the majority class, while Mazumder et al. (2025) argued that fraud rarity makes accuracy alone an unreliable performance measure.

Deep learning provides a stronger foundation for detecting nonlinear and temporal fraud patterns. Temporal Convolutional Networks are particularly suitable because dilated causal convolutions can learn short-term and long-range transaction dependencies while supporting efficient inference. Abd-Ellatif et al. (2025) reported that adaptive deep learning can improve real-time fraud classification, while Ayub et al. (2025) demonstrated the value of deep learning for banking fraud detection. However, financial institutions require models that are not only accurate but also explainable and operationally efficient. This study therefore proposes an adaptive TCN-based framework supported by SHAP explanations and evaluated against Logistic Regression, Random Forest, and XGBoost baselines.

Related Works

Machine Learning Fraud Detection

Machine learning has been widely applied to fraud detection because it can learn decision patterns from transaction records. Logistic Regression provides a simple baseline for estimating fraud probability, while Random Forest and boosting methods improve predictive stability by combining multiple weak learners. Khalid et al. (2024) found that ensemble learning improves credit card fraud detection performance, and Adamu-Fika et al. (2025) reported strong comparative performance for Random Forest and related machine learning models. Despite these strengths, static machine learning models may become less effective when fraud tactics change over time.

Deep Learning Approaches

Deep learning approaches have been introduced to capture complex and nonlinear fraud patterns. Yaganti (2023) applied autoencoder-driven anomaly detection, while Mazumder et al. (2025) proposed a unified deep learning approach for financial transaction fraud. Mohammad and Logeshwaran (2024) also showed that deep learning can improve real-time credit card fraud classification. Nevertheless, many deep learning models remain difficult to interpret, creating trust and regulatory challenges in financial decision-making environments.

TCN Applications

Temporal Convolutional Networks have gained attention because they can model sequence dependencies without the recurrent bottlenecks associated with RNN and LSTM models. Abd-Ellatif et al. (2025) showed that adaptive deep learning architectures can support real-time fraud detection, while Lawati et al. (2025) emphasized adaptive temporal processing for drift-aware fraud monitoring. Graph-based and reinforcement learning approaches have also been explored by Zakaria et al. (2025), Cui et al. (2025), and Cao et al. (2025), but these models may be computationally demanding for low-latency transaction authorization.

Imbalanced Data Techniques

Imbalanced data remains a central challenge in fraud detection because fraudulent transactions are rare. Patel et al. (2025) demonstrated that hybrid balancing methods such as SMOTE-ENN can improve minority fraud detection, while Adejoh et al. (2025) argued that adaptive learning can improve fraud recognition under skewed distributions. However, oversampling can introduce synthetic bias, and undersampling can remove useful legitimate transaction information. This paper therefore combines imbalance-aware evaluation with comparative modelling and temporal validation rather than relying on accuracy alone.

MATERIALS AND METHODS

Dataset Description

The study used the Kaggle European credit card fraud dataset, which contains anonymized transaction records with PCA-transformed features V1 to V28, transaction time, transaction amount, and a binary class label. The dataset contains 284,807 transactions, including 492 fraudulent transactions and 284,315 legitimate transactions. Fraud accounts for only 0.173 percent of the dataset, giving an imbalance ratio of approximately 577.88:1.

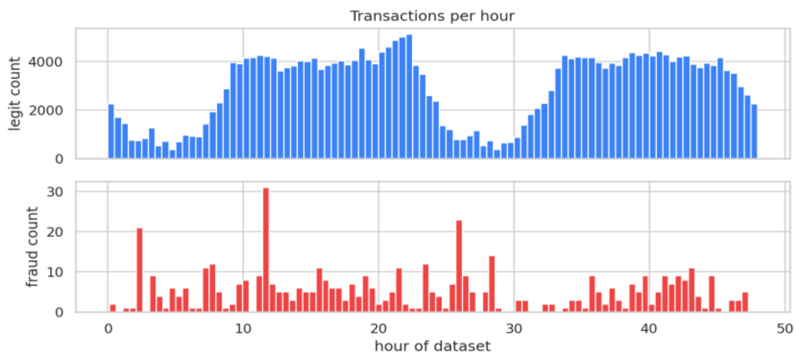

Figure 2: Transaction Distribution over Time VB

Table 1: Dataset Characteristics and Class Distribution

Characteristic	Value
Total transactions	284,807
Fraudulent transactions	492 (0.173%)
Legitimate transactions	284,315 (99.827%)
Class imbalance ratio	577.88:1
Temporal span	172,792 seconds, approximately 48 hours
Number of features	30: V1-V28, Time, Amount
Missing values	0
Duplicate transactions	1,081 (0.38%)

Data Pre-processing

Preprocessing transformed the raw transaction data into a structured format suitable for machine learning and deep learning models. The process included duplicate checking, missing value verification, feature scaling, and normalization. Since transaction features exist on different numerical scales, standardization was applied to improve training stability and prevent features with larger values from dominating model learning.

Feature Engineering

Feature engineering focused on transforming transaction attributes into meaningful representations for fraud classification. The PCA-transformed features were retained because they preserve anonymized transaction structure while protecting sensitive financial information. The transaction amount and time features were standardized and used alongside V1 to V28. The engineered feature vector supported Logistic Regression, Random Forest, XGBoost, and TCN training.

Temporal Split Strategy

The dataset was split chronologically rather than randomly to prevent future information leakage and simulate real-world deployment. In a production fraud detection system, models are trained on historical transactions and tested on future transactions. The chronological split also revealed concept drift, because fraud prevalence declined from 0.193 percent in training to 0.122 percent in testing. This 36.8 percent reduction indicates that fraud behaviour changes over time, supporting the need for adaptive learning (Lawati et al., 2025; Cao et al., 2025).

Table 2: Temporal Split Summary

Split	Size	Fraud count	Fraud rate	Time range
Training	199,364	384	0.193%	0-132928 seconds
Validation	42,721	56	0.131%	132929-151328 seconds
Test	42,722	52	0.122%	151328-172792 seconds

TCN Architecture

The proposed TCN model was designed for sequential fraud detection. It used four convolutional blocks with dilation factors of 1, 2, 4, and 8, allowing the network to capture both short-range and long-range transaction dependencies. Dropout was included to reduce overfitting, and the final sigmoid output generated the fraud probability score.

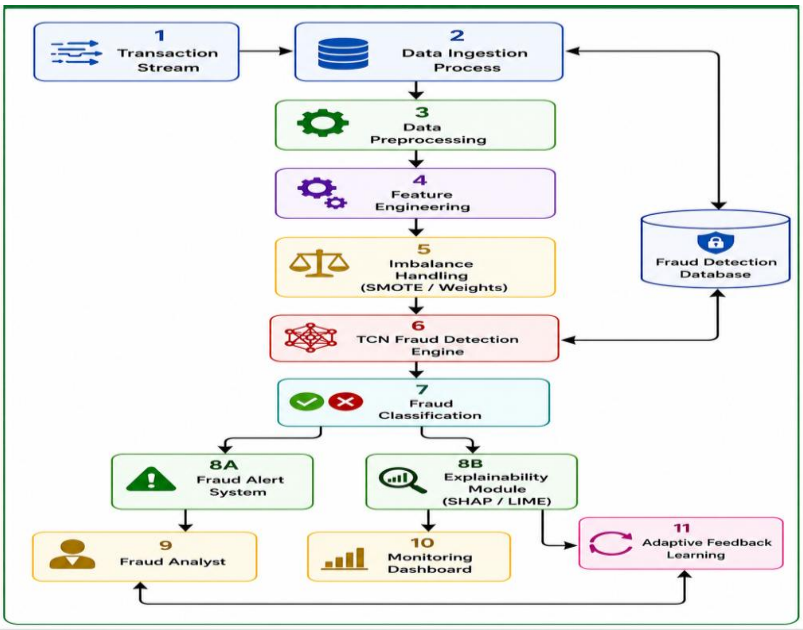

Figure 3: The proposed TCN model architecture

Table 3: TCN Architecture Specification

Component	Configuration
Input dimension	29 features
TCN blocks	4 blocks with dilations [1, 2, 4, 8]
Channels	[32, 32, 64, 64]
Kernel size	3
Dropout rate	0.2
Total parameters	54,945
Output layer	Fully connected 64 to 1 plus sigmoid

Hyperparameter Optimization

Hyperparameter optimization adjusted learning rate, convolution layers, kernel size, batch size, number of epochs, dropout rate, activation function, and optimizer configuration. The TCN was trained for 50 epochs with early stopping using patience of 10 epochs. The best validation PR-AUC of 0.721 was achieved at epoch 30, after which validation performance plateaued while training performance continued to improve.

RESULTS AND DISCUSSION

Training Dynamics

The TCN training curves showed rapid initial learning followed by gradual refinement. Early stopping prevented excessive overfitting after the validation PR-AUC stopped improving. The model maintained competitive performance despite using flattened transaction features rather than complete customer-level transaction sequences.

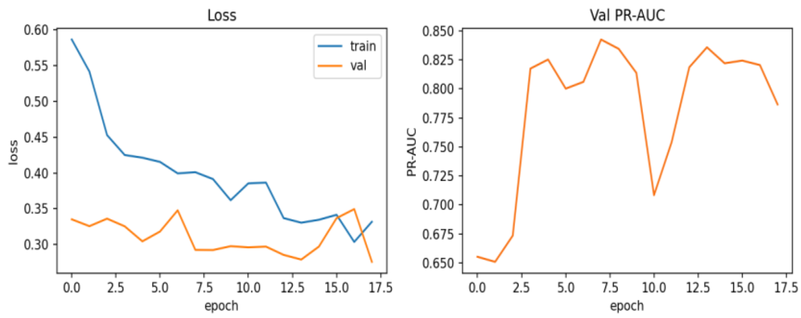

Figure 4: TCN Training and Validation Curves

ROC-AUC

XGBoost and Logistic Regression achieved the highest ROC-AUC values of 0.977, while TCN achieved a close ROC-AUC of 0.974. Random Forest recorded ROC-AUC of 0.959. These values indicate strong discrimination between fraudulent and legitimate transactions, although ROC-AUC must be interpreted alongside PR-AUC because the dataset is severely imbalanced.

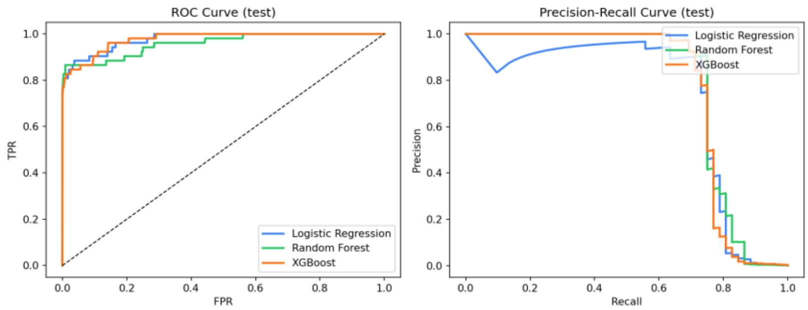

Figure 5: ROC and Precision-Recall Curves for All Models

Precision-Recall Analysis

Precision-recall performance provided a more realistic assessment of fraud detection under class imbalance. Random Forest achieved the highest PR-AUC of 0.773 and precision of 0.925, making it strongest for minimizing false alerts. XGBoost achieved PR-AUC of 0.759 and precision of 0.792. The TCN achieved PR-AUC of 0.714 and recall of 0.808, showing strong fraud sensitivity but lower precision at the default threshold.

Table 4: Comprehensive Model Performance on the Test Set

Model	Accuracy	PR-AUC	ROC-AUC	Precision	Recall	F1-score
Random Forest	99.96%	0.773	0.959	0.925	0.712	0.804
XGBoost	99.94%	0.759	0.977	0.792	0.731	0.760
TCN	99.06%	0.714	0.974	0.097	0.808	0.172
Logistic Regression	97.37%	0.711	0.977	0.038	0.846	0.073

Confusion Matrices

At the default classification threshold of 0.5, Random Forest generated only 3 false positives and detected 37 of 52 fraud cases. XGBoost generated 10 false positives and detected 38 fraud cases. TCN detected 42 fraud cases, giving the highest true positive rate of 80.8 percent, but it also generated 393 false positives. This indicates that the TCN requires threshold calibration before deployment in environments where false alerts are costly.

Table 5: Confusion Matrix Summary at Default Threshold

Model	False positives	True positives	False positive rate	True positive rate
Random Forest	3	37	0.007%	71.2%
XGBoost	10	38	0.023%	73.1%
TCN	393	42	0.92%	80.8%

Latency Evaluation

All models achieved sub-100 ms p99 latency, which is suitable for real-time transaction authorization. XGBoost and TCN were especially efficient, with p99 latency values of 6.40 ms and 6.34 ms respectively. Logistic Regression achieved the fastest inference, while Random Forest remained acceptable despite higher latency.

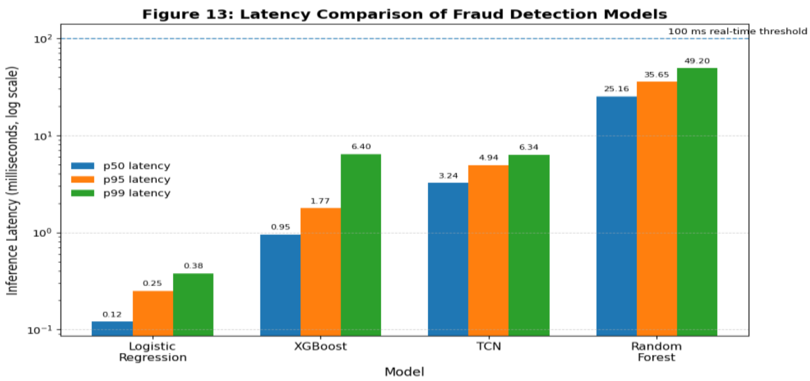

Figure 6: Latency Comparison of Fraud Detection

Table 6: Inference Latency Benchmarks

Model	p50 ms	p95 ms	p99 ms	Mean ms	Real-time viability
Logistic Regression	0.12	0.25	0.38	0.14	Excellent
XGBoost	0.95	1.77	6.40	1.06	Good
TCN	3.24	4.94	6.34	3.36	Good
Random Forest	25.16	35.65	49.20	26.35	Acceptable

Discussion

TCN Effectiveness

The TCN achieved competitive performance with PR-AUC of 0.714, ROC-AUC of 0.974, and recall of 0.808. Although it did not outperform Random Forest or XGBoost in precision, it showed strong ability to detect fraudulent transactions. This is important because fraud detection systems often prioritize identifying rare fraud cases before financial loss occurs. The lower precision suggests that the TCN should be calibrated and trained on richer sequential histories to improve operational precision.

Real-time Deployment Implications

The latency results demonstrate that the proposed framework can support real-time deployment. XGBoost and TCN both achieved p99 latency below 7 ms, leaving sufficient time for payment authorization workflows. Throughput testing in the thesis showed approximately 8,300 transactions per second for Logistic Regression, 650 for XGBoost, 620 for TCN, and 190 for Random Forest. These results indicate that horizontal scaling can support medium-sized financial institutions during peak transaction loads.

Comparison with Baseline Models

Random Forest was the best model for minimizing false positives, while XGBoost provided the best balance between predictive discrimination and deployment efficiency. TCN offered the strongest fraud sensitivity among the advanced models but required threshold tuning to reduce false positives. This confirms that no single model is universally superior; instead, model choice should depend on operational priorities such as fraud sensitivity, false alert cost, latency, interpretability, and scalability.

Explainability and Analyst Trust

SHAP analysis strengthened model transparency by identifying influential features behind fraud predictions. For Random Forest, V24 and V23 were the most influential variables, followed by V28, V27, V22, and V21. For XGBoost and TCN, V14 and V12 were especially influential. Although PCA anonymization prevents direct business interpretation of these variables, SHAP still provides feature-level explanations that can support analyst review, regulatory compliance, and model validation (Baisholan et al., 2025; Alessi and Fugini, 2026).

Table 1: SHAP Feature Importance Ranking for Transaction Fraud Detection

Rank	Feature	Mean absolute SHAP value
1	V24	0.0640
2	V23	0.0640
3	V28	0.0534
4	V27	0.0534
5	V22	0.0441
6	V21	0.0441
7	V8	0.0392
8	V7	0.0392
9	V20	0.0374
10	V19	0.0374

CONCLUSION

This paper presented an adaptive TCN-based framework for real-time credit card fraud detection in highly imbalanced financial transaction data. The study addressed digital fraud growth, limitations of conventional detection systems, severe class imbalance, temporal drift, real-time latency, and explainability. Experiments showed that Random Forest achieved the strongest PR-AUC and precision, XGBoost achieved the best overall balance between discrimination and latency, and TCN achieved strong fraud recall with real-time inference speed.

The findings show that ensemble models remain highly effective for structured credit card fraud datasets, while TCN provides a promising foundation for adaptive sequence-aware fraud detection. Future work should incorporate customer-level transaction sequences, optimize TCN thresholds, evaluate live streaming deployment, and integrate continuous learning for concept drift adaptation.

REFERENCES

Abd-Ellatif, L., Abrar, M., and Ismaeel, A. (2025). ATAD-Net: An Adaptive Deep Learning Framework for Real-Time Financial Fraud Detection. Adv. Artif. Intell. Mach. Learn., 5, 3988-4003. https://doi.org/10.54364/aaiml.2025.52225.

Adamu-Fika, F., Fatika, A., Baba-Onoja, A., Adeniyi, U., Mafua, H., Okpoko, O., and Ramalan, A. (2025). A Comparative Evaluation of Machine Learning Models for an Enhanced Fraud Detection in Financial Systems. Advances in Multidisciplinary & Scientific Research Journal Publication. https://doi.org/10.22624/aims/cisdi/v16n2p2.

Adejoh, J., Owoh, N., Ashawa, M., Hosseinzadeh, S., Shahrabi, A., and Mohamed, S. (2025). An Adaptive Unsupervised Learning Approach for Credit Card Fraud Detection. Big Data Cogn. Comput., 9, 217. https://doi.org/10.3390/bdcc9090217.

Al-Dahasi, E., Alsheikh, R., Khan, F., and Jeon, G. (2024). Optimizing fraud detection in financial transactions with machine learning and imbalance mitigation. Expert Systems, 42. https://doi.org/10.1111/exsy.13682.

Alessi, G., and Fugini, M. (2026). Adaptive Real-Time Financial Fraud Detection with Explainable AI tools. Digital Threats: Research and Practice. https://doi.org/10.1145/3794859.

Ayub, M., Bhattacharjee, B., Akter, P., Uddin, M., Gharami, A., Islam, M., Suhan, S., Khan, M., and Chambugong, L. (2025). Deep Learning for Real-Time Fraud Detection: Enhancing Credit Card Security in Banking Systems. The American Journal of Engineering and Technology. https://doi.org/10.37547/tajet/volume07issue04-19

Baisholan, N., Dietz, J., Gnatyuk, S., Turdalyuly, M., Matson, E., and Baisholanova, K. (2025). FraudX AI: An Interpretable Machine Learning Framework for Credit Card Fraud Detection on Imbalanced Datasets. Comput., 14, 120. https://doi.org/10.3390/computers14040120

Cao, J., Zheng, W., Ge, Y., and Wang, J. (2025). DriftShield: Autonomous Fraud Detection via Actor-Critic Reinforcement Learning With Dynamic Feature Reweighting. IEEE Open Journal of the Computer Society, 6, 1166-1177. https://doi.org/10.1109/ojcs.2025.3587001

Cui, Y., Han, X., Chen, J., Zhang, X., Yang, J., and Zhang, X. (2025). FraudGNN-RL: A Graph Neural Network With Reinforcement Learning for Adaptive Financial Fraud Detection. IEEE Open Journal of the Computer Society, 6, 426-437. https://doi.org/10.1109/ojcs.2025.3543450

Ibrahim, N., Kamal, R., and Nipo, D. (2024). Society's Growing Preference for Cashless Transactions Over Cash. International Journal of Academic Research in Business and Social Sciences. https://doi.org/10.6007/ijarbss/v14-i9/22781

Khalid, A., Owoh, N., Uthmani, O., Ashawa, M., Osamor, J., and Adejoh, J. (2024). Enhancing Credit Card Fraud Detection: An Ensemble Machine Learning Approach. Big Data Cogn. Comput., 8, 6. https://doi.org/10.3390/bdcc8010006

Lawati, H., Zainal, A., Al-Rimy, B., Al-Azawi, M., Kassim, M., Almalki, S., and Alghamdi, T. (2025). An Integrated Preprocessing and Drift Detection Approach With Adaptive Windowing for Fraud Detection in Payment Systems. IEEE Access, 13, 92036-92056. https://doi.org/10.1109/access.2025.3569609.

Lakhotia, K. (2024). Revolutionizing Digital Payments and Transforming Global Economies: A Drive Promoting Cashless Market Scenarios. International Journal of Science and Research. https://doi.org/10.21275/sr24406144130.

Mazumder, M., Shourov, M., Rasul, I., Akter, S., and Miah, M. (2025). Fraud Detection in Financial Transactions: A Unified Deep Learning Approach. Journal of Economics, Finance and Accounting Studies. https://doi.org/10.32996/jefas.2025.7.2.16.

Mohammad, R., and Logeshwaran, J. (2024). Real-Time Credit Card Fraud Detection using deep Learning based framework. 2024 2nd International Conference on Advances in Computation, Communication and Information Technology, 1, 496-500. https://doi.org/10.1109/icaiccit64383.2024.10912330.

Pasupuleti, M. (2025). Deep Learning for Fraud Detection in Real-Time Transaction Networks. International Journal of Academic and Industrial Research Innovations. https://doi.org/10.62311/nesx/rphcr24.

Patel, H., Patel, M., and Savani, K. (2025). An Optimized XGBoost Framework for Real-Time Credit Card Fraud Detection: Addressing Class Imbalance with Hybrid SMOTE-ENN Resampling. International Journal of Scientific Research in Science and Technology. https://doi.org/10.32628/ijsrst25123122.

Taralkar, J. (2025). FinAI: Deep learning for real-time anomaly detection in financial transactions. World Journal of Advanced Engineering Technology and Sciences. https://doi.org/10.30574/wjaets.2025.15.2.0354.

Teker, S., Teker, D., and Orman, I. (2022). DIGITAL PAYMENT SYSTEMS: A FUTURE OUTLOOK. Pressacademia. https://doi.org/10.17261/pressacademia.2022.1613.

Yaganti, D. (2023). Unsupervised Deep Learning for Credit Card Fraud Detection: An Autoencoder-Driven Framework with Real-Time Dash Visualization Using Tensorflow 2.X. International Journal of Advanced Research in Science, Communication and Technology. https://doi.org/10.48175/ijarsct-11978t.

Zakaria, R., Rahman, M., Tazwar, M., Choudhury, H., Rahman, H., and Rafi, M. (2025). Detecting Financial Fraud in Real-Time Transactions Using Graph Neural Networks and Anomaly Detection Techniques. Journal of Economics, Finance and Accounting Studies. https://doi.org/10.32996/jefas.2025.7.6.1.