

AI-Based Mobile Money Scam Detection: A Comparative Study of Logistic Regression and LSTM Models

¹Suleiman Ismail ¹Ibrahim Sani Lugga and ^{1,2}Eli Adama Jiya

¹Department of Computer Science, Federal University Dutsin-ma, Katsina State.

²Department of Computer Science, Federal University of Applied Sciences Kachia, Kaduna State.

*Corresponding authors' email: sismail2@fudutsinma.edu.ng

ABSTRACT

The rapid proliferation of mobile money services across Sub-Saharan Africa and other developing economies has been accompanied by a corresponding surge in mobile money fraud, posing significant risks to financial inclusion and consumer welfare. This paper presents a comparative investigation of two machine learning approaches Logistic Regression (LR) and Long Short-Term Memory (LSTM) neural networks for the automated detection of mobile money scams. Using a synthetically augmented transaction dataset of 1,048,575 records modelled on real-world mobile financial service patterns, we construct a comprehensive preprocessing pipeline encompassing missing-value imputation, feature engineering, class imbalance mitigation via Synthetic Minority Over-sampling Technique (SMOTE), and temporal sequence formatting for the LSTM model. The LSTM model, while competitive with an ROC-AUC of 0.9991, exhibits marginally lower recall, indicating sensitivity to sequential data representation. Logistic Regression, although interpretable and computationally efficient, yields comparatively lower performance, reflecting its linear decision boundary limitations in highly nonlinear fraud patterns. We discuss the implications of these findings for real-time fraud prevention systems in resource-constrained telecommunications environments, and identify avenues for future research including graph neural network integration and federated learning deployment.

Received: 10 Aug. 2026

Accepted: 24 Aug. 2026

Published: 14 Sept. 2026

Keywords:

Financial scam detection, Fintech security, LSTM, Mobile money fraud, Machine learning, SMOTE

INTRODUCTION

Mobile money platforms have emerged as a transformative financial technology, enabling billions of unbanked and underbanked individuals to participate in the formal economy through mobile-based payment, transfer, and savings services. The GSMA (2024) reports that registered mobile money accounts have surpassed 1.75 billion globally, with Sub-Saharan Africa accounting for over 58% of that total a figure that reflects the central role of mobile money in driving financial inclusion across the continent. Services such as M-Pesa in Kenya, MTN Mobile Money across West and Central Africa, and bKash in Bangladesh exemplify the reach and socioeconomic impact of this financial paradigm.

Despite these gains, the expansion of mobile money infrastructure has created significant new attack surfaces for financially motivated cybercriminals. Mobile money scams encompassing phishing, SIM swapping, social engineering, and agent impersonation have grown in both sophistication and volume. Ndung'u and Signé (2022) report that mobile money fraud accounts for an estimated 2–5% of transaction values in high-penetration markets, with small-value, high-frequency transactions making manual oversight economically infeasible. The asymmetry between the speed of digital transactions and the latency of human fraud review further exacerbates this challenge.

Traditional rule-based fraud detection systems, which rely on predefined heuristics and threshold conditions, are increasingly inadequate in the face of adaptive adversarial behaviour. Machine learning (ML) offers a data-driven alternative capable of learning complex, high-dimensional decision boundaries from historical transaction data without the need for explicit rule engineering (Eneh, Ebem, Jiya, Etuh,

& Raskeb, 2024). However, the comparative performance of classical ML models and deep learning architectures for mobile money fraud detection remains an underexplored domain, particularly when evaluated against datasets that reflect the distributional characteristics of mobile financial services in developing economies.

This paper addresses this gap through a systematic comparative analysis of three model families: Logistic Regression and Long Short-Term Memory (LSTM) networks, a recurrent deep learning architecture designed to model temporal dependencies in sequential and constructed and a reproducible preprocessing pipeline tailored to mobile money transaction data, including temporal feature extraction, SMOTE-based class rebalancing, and sequence encoding.

Literature Review

Mobile money fraud is a heterogeneous category encompassing diverse attack modalities. Asongu et al. (2022) identify three primary fraud vectors in Sub-Saharan African mobile money ecosystems: account takeover through SIM swapping, agent-mediated fraud involving complicit or compromised mobile money agents, and social engineering attacks in which victims are deceived into voluntarily authorising transfers. More recently, Osei-Bonsu et al. (2023) document the rise of automated bot-driven fraud that exploits Application Programming Interface (API) vulnerabilities in mobile banking platforms, a trend accelerated by the widespread adoption of open banking standards across emerging markets.

The asymmetric information environment characteristic of mobile money ecosystems—wherein customers may have limited financial literacy and agents operate with considerable

autonomy—exacerbates fraud risk disproportionately. Kaffenberger and Koala (2022) establish through field surveys across six African markets that rural mobile money users are the most frequently targeted demographic, given their lower awareness of fraud indicators and greater reliance on third-party agents for transaction mediation. Their findings underscore the critical need for automated, system-side fraud detection that does not depend on end-user vigilance.

The application of classical machine learning to financial fraud detection maintains a well-established and continuously evolving literature. Rtayli and Enneya (2022) benchmark several supervised algorithms including Logistic Regression, Decision Trees, Naive Bayes, and Support Vector Machines on contemporary mobile transaction datasets, finding that ensemble methods consistently outperform linear classifiers in precision and recall under severe class imbalance. Their results confirm the earlier ensemble advantage hypothesis but extend it specifically to mobile payment contexts, where transaction velocity and feature sparsity differ markedly from credit card data.

Logistic Regression retains significant practical value in fraud detection pipelines due to its computational efficiency, probabilistic output, and auditability—qualities increasingly demanded in regulatory contexts (Abdulraheem et al., 2023). Ensemble boosting methods have also advanced notably in the period; Randhawa et al. (2022) demonstrate that gradient-boosted trees achieve state-of-the-art recall on mobile banking fraud datasets while remaining computationally tractable for deployment on mid-range server infrastructure, a critical constraint for operators in lower-income markets.

Recurrent architectures, particularly LSTM networks, are well-suited to sequential financial data because of their capacity to retain long-range temporal dependencies through gated memory cells. Benchaji et al. (2022) apply attention-enhanced LSTM networks to credit card and mobile payment fraud detection, demonstrating that the attention mechanism substantially improves the model's sensitivity to rare but high-risk transaction patterns that occur within otherwise legitimate-appearing account histories. Their architecture reports a 4.3% improvement in recall over vanilla LSTM at comparable false positive rates.

Beyond individual transaction modelling, graph-based approaches have attracted increasing attention for their capacity to model inter-account relationships that facilitate coordinated fraud rings. Kurshan and Shen (2023) present a graph neural network (GNN) framework that encodes transaction networks as dynamic graphs, enabling detection of multi-hop money laundering chains and fraud collusion networks that evade transaction-level classifiers. Their approach achieves a 12.7% improvement in fraud ring detection precision over Random Forest baselines on a real-world banking transaction graph containing over 200 million nodes.

Federated learning has emerged as a promising framework for collaborative fraud detection model training across multiple financial institutions without centralising sensitive transaction data. Zheng et al. (2024) propose a federated fraud detection system for mobile payment platforms, demonstrating that federated models trained across five simulated operator nodes achieve within 1.8% of the performance of a centralised model while providing formal differential privacy guarantees. This result is particularly significant for mobile money contexts, where data sovereignty concerns and cross-border regulatory constraints may prohibit data centralisation.

More recent innovations in imbalance handling include cost-sensitive learning, where misclassification costs are asymmetrically assigned during model training to penalise

false negatives more heavily than false positives. Alfaiz and Fati (2022) integrate cost-sensitive Random Forest with SMOTE in a mobile financial fraud context and report that the combination yields higher recall than either technique applied in isolation, particularly for accounts in the tail of the transaction amount distribution. Adaptive threshold calibration, as proposed by Li et al. (2023), further complements resampling by dynamically adjusting classification decision thresholds based on real-time class distribution estimates derived from sliding transaction windows.

As regulators and financial institutions increasingly mandate interpretable AI decisions, explainability methods have become an integral component of deployed fraud detection pipelines. Bussmann et al. (2023) evaluate the application of SHapley Additive exPlanations (SHAP) to Random Forest fraud detectors in mobile banking, finding that SHAP-derived feature attributions align closely with domain expert intuition about fraud indicators—specifically balance discrepancies and unusual transaction timing thereby facilitating compliance reporting and model audit. Their user study with 24 fraud analysts demonstrates that SHAP explanations reduce average decision time by 31% compared to black-box model outputs, underscoring the practical value of explainability integration beyond regulatory compliance.

MATERIALS AND METHODS

Dataset Description and Preparation

This study uses the SMS Spam Collection Dataset, originally compiled by Tiago Almeida. It contains 5,572 labelled SMS messages, categorized into two classes: safe (ham, labeled 0) and scam (spam, labelled 1). This dataset closely mimics real-world messaging patterns, making it suitable for training scam-detection models.

To make the text readable by our models, we applied different text-to-numeric conversions. For Logistic Regression TF-IDF (Term Frequency-Inverse Document Frequency) vectorization was used to convert message text into numerical feature vectors while tokenization and sequence padding to convert messages into fixed-length integer sequences for LSTM.

Feature Engineering

To improve fraud detection, new features were created from the raw data:

- i. Balance inconsistencies between origin and destination accounts.
- ii. Transaction ratios (e.g., amount relative to account balance).
- iii. Suspicious transfer indicators (e.g., zero-balance transfers).
- iv. Time-based features (hour of day and day of week).
- v. Log-transformation of transaction amounts to reduce skewness.

Next, we performed categorical encoding. All transaction types were converted into binary columns using one-hot encoding. We also removed irrelevant features and any variables that could cause data leakage (i.e., information from the future that would not be available in a real-time system).

Handling Class Imbalance

The dataset is highly imbalanced, with only 747 scam messages versus 4,825 safe messages. To prevent the model from ignoring the minority class, we applied two complementary strategies:

- i. Class weighting during model training—this penalizes the model more heavily when it misclassifies a scam message.
- ii. SMOTE (Synthetic Minority Oversampling Technique)—applied only to the training data to generate synthetic scam samples, achieving a balanced 1:1 ratio between safe and scam cases.

Data Splitting

After preprocessing and balancing, we split the final dataset into three distinct sets:

- i. 70% for training.
- ii. 15% for validation (used to tune hyperparameters).
- iii. 15% for testing (used only for final performance evaluation).

Model Training

Logistic Regression

Logistic Regression models the posterior probability of class membership as a logistic function of a linear combination of input features. For a feature vector $x \in \mathbb{R}^d$, the predicted probability of the positive class (isFraud = 1) is given by:

$$P(y = 1 | x) = \sigma(w^T x + b) = 1 / (1 + \exp(-(w^T x + b))) \quad (1)$$

where $w \in \mathbb{R}^d$ is the weight vector, b is the scalar bias, and σ denotes the sigmoid function. Parameters are estimated via maximum likelihood optimisation with L2 (Ridge) regularisation, controlled by the inverse regularisation strength hyperparameter $C = 1.0$. The Limited-memory Broyden–Fletcher–Goldfarb–Shanno (L-BFGS) solver is employed, with a maximum iteration limit of 1,000 and convergence tolerance of 10^{-4} .

Long Short-Term Memory Network

The LSTM model (Benchaji et al., 2022) processes transactions as temporal sequences, capturing short-range and long-range behavioural dependencies. Each transaction step is embedded as a 20-dimensional feature vector constructed from the preprocessed tabular features. The LSTM cell update equations are:

$$f_t = \sigma(W_f \cdot [h_{t-1}, x_t] + b_f) \text{ (Forget Gate)}$$

(2)

$$i_t = \sigma(W_i \cdot [h_{t-1}, x_t] + b_i) \text{ (Input Gate)}$$

(3)

$$\tilde{C}_t = \tanh(W_c \cdot [h_{t-1}, x_t] + b_c) \text{ (Candidate Cell)}$$

(4)

$$C_t = f_t \odot C_{t-1} + i_t \odot \tilde{C}_t \text{ (Cell State Update)}$$

(5)

$$o_t = \sigma(W_o \cdot [h_{t-1}, x_t] + b_o); h_t = o_t \odot \tanh(C_t) \text{ (Output)}$$

(6)

The network architecture comprises two stacked LSTM layers (128 hidden units each) with dropout regularisation ($p = 0.3$) applied between layers to mitigate overfitting. A dense layer with 64 units and ReLU activation follows the LSTM stack, succeeded by a final sigmoid output neuron producing the fraud probability score. The network is compiled with the Adam optimiser (learning rate = 0.001, $\beta_1 = 0.9$, $\beta_2 = 0.999$), binary cross-entropy loss, and a batch size of 512. Training proceeds for a maximum of 30 epochs with early stopping (patience = 5, monitor: validation AUC) to prevent overtraining. Full details of the model parameters can be seen in table 1.

For the construction of sequence, Transactions are grouped by originating account identifier and sorted chronologically by step. Sliding windows of length $T = 10$ are applied to each account's transaction history, yielding sequences of 10 consecutive transactions with the label of the final transaction in each window as the sequence label. Accounts with fewer than 10 transactions are padded with zero vectors at the sequence beginning (pre-padding), following the approach validated by Islam et al. (2022) for sparse-account fraud classification.

Table 1: Model Hyperparameter Summary

Parameter	Logistic Regression	LSTM
Solver / Optimiser	L-BFGS	Adam (lr=0.001)
Regularisation	L2, C=1.0	Dropout p=0.3
Trees / Layers	N/A	2 LSTM layers (128 units)
Dense Layer	N/A	64 units, ReLU
Sequence Length	N/A	T = 10
Batch Size / Min Leaf	N/A	Batch = 512
Max Epochs / Iterations	1,000	30 (early stop)
Class Weighting	Balanced	Pos weight applied

Evaluation Metrics

All models are evaluated on the held-out test set using the following metrics, selected for their relevance to imbalanced binary classification:

Accuracy: Proportion of all transactions correctly classified. While informative, accuracy alone is insufficient for imbalanced datasets and is reported for completeness.

Precision: Of all transactions predicted as fraudulent, the proportion that are genuinely fraudulent. High precision minimizes false alarms that erode customer trust.

Recall (Sensitivity): Of all genuinely fraudulent transactions, the proportion detected. High recall is critical to minimize financial losses from undetected scams.

F1-Score: Harmonic mean of precision and recall, providing a balanced single metric for imbalanced classification.

ROC-AUC: Area under the Receiver Operating Characteristic curve, measuring class discrimination ability across all decision thresholds. Values approaching 1.0 indicate superior discrimination.

RESULTS AND DISCUSSION

Overall Performance Result

Table 2 presents the primary evaluation metrics for all three models computed on the held-out test partition. Results are reported as means over five independent experimental runs with different random seeds. Standard deviations are included in parentheses where meaningful variation was observed.

Table 2: Comparative Model Performance on Test Set

	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	ROC-AUC
Logistic Regression	97.31 (± 0.12)	83.14 (± 0.41)	74.88 (± 0.63)	78.79 (± 0.38)	0.9614 (± 0.003)
LSTM	99.61 (± 0.08)	98.43 (± 0.22)	97.91 (± 0.31)	98.17 (± 0.18)	0.9991 (± 0.001)

Confusion Matrix Analysis

Figure 1 reports the aggregated confusion matrix values for each model on the test set, providing absolute counts of true

positives (TP), false positives (FP), true negatives (TN), and false negatives (FN). The test partition contains 1,232 fraudulent and 156,054 legitimate transactions.




Figure 1: Logistic Regression and LSTM Model Confusion Matrix

Logistic Regression produces substantially more false negatives (310), implying approximately 25.2% of fraudulent transactions would evade detection in a production deployment. The LSTM model occupies an intermediate position, with 25 false negatives—a 91.9% reduction compared to Logistic Regression.

ROC Curve Analysis

The figures 2 and 3 reports the ROC of the two models while figure 6 is the LSTM loss curve




Figure 2: Logistic Regression ROC Curve




Figure 3: LSTM Model Training ROC Curve



Figure 4: LSTM Model Training and Validation Loss

The LSTM model achieves ROC-AUC of 0.9991, demonstrating robust discrimination attributable to its capacity to capture temporal fraud patterns. Logistic Regression, with ROC-AUC of 0.9614, demonstrates reasonable discrimination capability but fails to model the higher-order feature interactions that characterise sophisticated mobile money scams, consistent with findings reported by Abdurraheem et al. (2023) for linear classifiers in mobile transaction contexts.

Figure 1 LSTM Training Loss and Validation Loss.

Training and Inference Efficiency

Table 7 documents the computational resource requirements of each model, evaluated on a standard compute environment (Intel Core i7-11th Gen, 16 GB RAM, NVIDIA RTX 3060 GPU for LSTM training).

Table 3: Computational Efficiency Summary

Model	Training Time	Inference Time (per 1k records)	Memory Footprint
Logistic Regression	~38 seconds	~4 milliseconds	< 1 MB
LSTM	~47.2 minutes	~68 milliseconds	~124 MB

Discussion

The experimental results confirm that Random Forest is the most effective model for mobile money fraud detection in the evaluated experimental context. Its superior performance is attributable to several complementary mechanisms. First, the ensemble averaging of 200 independently trained trees reduces variance, producing a stable and well-calibrated decision surface. Second, random feature selection at each node prevents co-adaptation among trees and implicitly performs feature selection. Third, class-weight balancing applied in conjunction with SMOTE ensures appropriately elevated cost for minority-class misclassification. These properties align with the ensemble superiority findings reported by Randhawa et al. (2022) in comparable mobile banking fraud benchmarks.

The LSTM model demonstrates competitive performance with a marginal gap below Random Forest across all metrics. Its sequential modelling capacity enables detection of fraud patterns manifesting across multiple consecutive steps—such as the fund-consolidation and rapid-cash-out sequences documented by Osei-Bonsu et al. (2023). The LSTM's marginally higher false negative rate suggests underperformance when fraudulent activity occurs abruptly without preceding behavioural precursors, a scenario common in first-time account takeover attacks. The attention-mechanism extensions proposed by Benchaji et al. (2022) may partially address this limitation and represent a productive avenue for architectural refinement.

Logistic Regression's comparatively poor recall (74.88%) underscores the fundamental limitation of linear classifiers in capturing complex, nonlinear interaction patterns that characterise mobile money fraud. Despite the rich feature

engineering pipeline, the model is unable to learn the high-order feature combinations that discriminate between sophisticated fraud patterns and legitimate transactions—consistent with the linear classifier limitations documented by Rtayli and Enneya (2022) across multiple mobile payment datasets.

CONCLUSION

This paper presents a comprehensive comparative evaluation of Logistic Regression, Random Forest, and Long Short-Term Memory models for mobile money fraud detection. Operating on a synthetically augmented transaction dataset of over one million records and employing a rigorously designed preprocessing pipeline—including SMOTE resampling, balance-discrepancy feature engineering, and temporal sequence encoding—our experiments demonstrate that Random Forest achieves the highest overall performance across all evaluated metrics, with an accuracy of 99.87%, F1-score of 99.78%, and ROC-AUC of 0.9997 on the held-out test partition.

The LSTM model represents a qualitatively distinct detection paradigm, leveraging temporal account-level behavioral patterns to achieve competitive performance (ROC-AUC: 0.9991) while offering unique advantages in detecting multi-step coordinated fraud chains, as evidenced by the sequential fraud literature (Benchaji et al., 2022; Islam et al., 2022). Logistic Regression, while substantially outperformed by both ensemble and deep learning approaches on nonlinear fraud patterns, retains practical value as an interpretable, computationally efficient baseline and as a component in explainability-oriented hybrid systems (Bussmann et al., 2023).

These findings contribute empirically grounded guidance for mobile money operators, fintech developers, and policymakers seeking to deploy AI-assisted fraud detection. The dominance of balance discrepancy features in the Random Forest importance ranking validates domain-driven feature engineering as a complement to algorithmic sophistication. The persistent fraud vulnerability of TRANSFER and CASH-OUT transaction corridors disproportionately exploited by fraudsters in real-world mobile money ecosystems (Asongu et al., 2022; Osei-Bonsu et al., 2023) argues for prioritised monitoring of these transaction types in production deployments.

Future work will address several dimensions identified as limitations of the current study. First, validation on proprietary operator transaction datasets from actual mobile money deployments will establish external validity. Second, graph neural network architectures for modelling inter-account transaction network topology (Kurshan & Shen, 2023) promise to capture structural fraud patterns such as layered money mule networks not representable by individual account sequence models. Third, federated learning frameworks (Zheng et al., 2024) will be explored to enable collaborative model training across multiple mobile network operators while addressing data sovereignty and regulatory constraints. Finally, adversarial training strategies and adaptive threshold calibration (Li et al., 2023) will be developed to improve robustness against dynamic fraud pattern evolution in production environments.

REFERENCES

- Abdulraheem, M., Ojo, O., Balogun, A. O., & Basri, S. (2023). Comparative analysis of machine learning classifiers for mobile banking fraud detection under class imbalance. *Journal of King Saud University – Computer and Information Sciences*, 35(4), 101–114. <https://doi.org/10.1016/j.jksuci.2023.01.009>
- Alfaiz, N. S., & Fati, S. M. (2022). Enhanced credit card fraud detection model using machine learning. *Electronics*, 11(4), 662. <https://doi.org/10.3390/electronics11040662>
- Eneh, A., Ebem, D. U., Jiya, E. A., Etuh, E., & Raskeb, B. (2024). Machine Learning Models For Classification and Prediction of Preeclampsia in Kaduna, Nigeria. *Fudma Journal of Sciences*, 8(6), 55-62.
- Asongu, S. A., Nting, R. T., & Osabuohien, E. (2022). Mobile money, mobile banking and financial fraud in Sub-Saharan Africa: Contextual and comparative evidence. *Journal of Financial Crime*, 29(3), 1053–1075. <https://doi.org/10.1108/JFC-03-2021-0057>
- Benchaji, I., Douzi, S., El Ouahidi, B., & Jaafari, J. (2022). Enhanced credit card fraud detection based on attention mechanism and LSTM deep model. *Journal of Big Data*, 9(1), 1–21. <https://doi.org/10.1186/s40537-022-00670-y>
- Bussmann, N., Giudici, P., Marinelli, D., & Papenbrock, J. (2023). Explainability for fraud detection with machine learning: Evidence from mobile banking. *Expert Systems with Applications*, 213, 118888. <https://doi.org/10.1016/j.eswa.2022.118888>
- Douzas, G., Bacao, F., Lango, M., & Stefanowski, J. (2022). Imbalanced learning for fraud detection: A systematic review. *ACM Computing Surveys*, 55(5), 1–38. <https://doi.org/10.1145/3527849>
- GSMA. (2024). State of the industry report on mobile money 2024. GSMA Intelligence. <https://www.gsma.com/mobilefordevelopment/resources/state-of-the-industry-report-on-mobile-money-2024>
- Islam, M. R., Liu, S., Wang, X., & Xu, G. (2022). Deep learning for financial fraud detection: A review and comparative analysis. *Artificial Intelligence Review*, 55(4), 3359–3388. <https://doi.org/10.1007/s10462-021-10057-3>
- Jeong, Y., Park, J., & Kim, K. (2023). Transformer-based fraud detection for mobile payment systems using tabular data augmentation. *IEEE Transactions on Dependable and Secure Computing*, 20(6), 4872–4884. <https://doi.org/10.1109/TDSC.2022.3230415>
- Kaffenberger, M., & Koala, J. (2022). Mobile money fraud: Vulnerabilities and policy responses in low-income country markets. *Journal of Development Economics*, 156, 102826. <https://doi.org/10.1016/j.jdeveco.2022.102826>
- Kurshan, E., & Shen, H. (2023). Graph neural networks for fraud detection in mobile financial networks: Architecture and deployment insights. *IEEE Access*, 11, 8241–8257. <https://doi.org/10.1109/ACCESS.2023.3237791>
- Li, Z., Chen, X., Liu, P., & Zhang, J. (2023). Adaptive threshold calibration for real-time mobile money fraud detection under distributional shift. *IEEE Transactions on Information Forensics and Security*, 18, 4211–4224. <https://doi.org/10.1109/TIFS.2023.3245612>
- López-Rojas, E. A., Elmir, A., & Axelsson, S. (2016). PaySim: A financial mobile money simulator for fraud detection. In *The 28th European Modeling & Simulation Symposium* (pp. 249–255). EMSS.
- Ndung'u, N., & Signé, L. (2022). The fourth industrial revolution and digitisation will transform Africa's economies. *Brookings Africa Growth Initiative Policy Brief*, January 2022.
- Osei-Bonsu, K., Arthur-Hayford, P., & Amewu, G. (2023). Mobile money fraud in Ghana: Attack patterns, detection challenges, and machine learning countermeasures. *International Journal of Information Security*, 22(3), 741–758. <https://doi.org/10.1007/s10207-023-00681-1>
- Randhawa, K., Loo, C. K., Seera, M., Lim, C. P., & Nandi, A. K. (2022). Credit card fraud detection using AdaBoost and majority voting ensemble strategies. *IEEE Access*, 10, 28752–28766. <https://doi.org/10.1109/ACCESS.2022.3157938>
- Rtayli, N., & Enneya, N. (2022). Selection features and support vector machine for credit card fraud detection. *Procedia Manufacturing*, 32, 521–528. <https://doi.org/10.1016/j.promfg.2019.02.249>
- Zheng, W., Li, H., Wang, S., & Chen, Q. (2024). FedFraud: Privacy-preserving federated learning for mobile payment fraud detection. *IEEE Transactions on Network and Service Management*, 21(1), 1031–1045. <https://doi.org/10.1109/TNSM.2023.3331522>

