FUDMA Journal of Engineering and Technology
ISSN: 3092-9385
https://fjet.fudutsinma.edu.ng/index.php/fjet/
Volume 2, Issue 2, 2026, 391-397

Detecting and Mitigating Cybersecurity Threats in Nigerian Commercial Banks Using Machine Learning: A Case Study of Selected Institutions

Folakemi B. OYEYEMI[1*], Usman A. ABDURRAHMAN[2]

[1*]Department of Information Technology, School of Computing, MIVA Open University, Abuja, Nigeria
[2]Department of Information and Communication Technology, Faculty of Computing, Northwest University, Kano, Nigeria

[1*]folakemi8982@gmail.com, [2]usmanalhaji79@gmil.com

Abstract

Nigerian commercial banks are facing rising cybersecurity threats, with electronic fraud losses surpassing ₦5.1 billion in 2023 [1], [2]. Traditional signature-based detection systems are reactive and struggle against dynamic, advanced attacks. This study developed and tested a machine learning-based cybersecurity threat detection and automated mitigation system designed for Nigerian commercial banks, with GTBank and Access Bank as institutional case studies [3], [4]. Guided by the CRISP-DM methodology [5], [6], [7], the study built and tested three classifiers: Random Forest, Support Vector Machine (SVM), and Long Short-Term Memory (LSTM), across two benchmark datasets, the CICIDS 2017 network intrusion dataset and the IEEE-CIS Financial Fraud Detection dataset. To correct pronounced class imbalance without data leakage into the test set, SMOTE oversampling was restricted to the training partition only. The results showed that all models exceeded the minimum macro F1-score of 0.90 for network intrusion detection. Random Forest achieved near-perfect classification (macro F1: 0.9974), followed by LSTM (0.9468) and SVM (0.9203). For financial fraud detection, which presented extreme class imbalance (3.5% fraud rate), Random Forest achieved the highest macro F1-score (0.7616), while LSTM demonstrated the strongest fraud recall (0.76), highlighting its suitability for sequential threat analysis. These results are consistent with recent studies on ML-based fraud detection using imbalanced datasets [8], [9], [10], [11]. An automated mitigation module was integrated via a Flask RESTful API, enabling real-time responses. While direct experimental comparison with a deployed signature-based IDS was outside the scope of this study, the results provide strong evidence that the evaluated ensemble and deep learning models substantially exceed performance thresholds associated with rule-based detection systems documented in the literature [12], [13]. The findings therefore suggest that ensemble and deep learning models offer a superior alternative to traditional signature-based systems, providing a high-accuracy threat detection and response framework for the Nigerian financial sector.

Keywords: *Machine Learning, Cybersecurity, Financial Fraud, Random Forest, LSTM.*

1.0 Introduction

The global financial services sector is undergoing a major digital transformation. In Nigeria, expanding digital-payment infrastructure, coupled with the CBN's drive toward broader financial inclusion, has pushed digital banking adoption forward at pace [14].

As a result, banks now move trillions of naira through digital channels every day [15]. That expanded digital footprint has turned Nigerian banks into attractive targets for advanced attacks: ransomware, DDoS campaigns, SIM-swap fraud, and identity theft built around compromised BVN records [16]. Worldwide, the annual cost of cybercrime is forecast to climb to $13.82 trillion by 2028 [17], and Nigerian financial institutions carry a disproportionate share of that burden given how quickly their services have digitized. The Nigeria Inter-Bank Settlement System (NIBSS) reported that electronic fraud losses in the sector exceeded ₦5.1 billion in 2023, exposing the inadequacy of current countermeasures [1].

Traditional security mechanisms, like signature-based intrusion detection systems (IDS) and static firewalls, depend on predefined rules to detect known threats. They are inherently reactive and have difficulty identifying new zero-day vulnerabilities or changing fraud patterns [18], [12], [19]. This often leads to high rates of false positives that overwhelm security operations teams. Nigeria's National Cybersecurity Policy and Strategy [20] identifies these limitations as a critical national security concern, urging the adoption of adaptive, intelligence-driven detection systems.

Machine learning (ML) offers a transformative alternative by analyzing large volumes of network and transaction data in real-time to identify anomalies and previously unseen attack patterns [21]. Supervised methods such as Random Forest and SVM, and deep-learning architectures such as LSTM, have already shown strong results in cybersecurity applications elsewhere [12], [13].

However, there is a lack of targeted research addressing specific threat vectors prevalent in the Nigerian banking environment [19]. Furthermore, existing studies rarely combine threat detection with automated real-time mitigation protocols.

Received: 16/06/2026; Revised: 25/07/2026; Accepted: 27/07/2026; Published: 01/09/2026

This study aims to bridge this gap by designing, implementing, and evaluating a machine learning-based cybersecurity threat detection and mitigation system tailored for Nigerian commercial banks, specifically GTBank and Access Bank, as institutional case studies. Nigerian commercial banking fraud scenarios that this study sought to address include: SIM swap attacks, where fraudsters hijack a customer's mobile number to intercept one-time passwords (OTPs) and authorize unauthorized fund transfers; Bank Verification Number (BVN)-linked identity theft, where compromised BVN data is exploited to open fraudulent accounts or obtain unauthorized loans; credential phishing leading to unauthorized internet and mobile banking transfers; and large-scale DDoS attacks targeting core banking infrastructure to disrupt services. These threat patterns, documented by [1], [22], and [2], directly informed the selection of the CICIDS 2017 and IEEE-CIS datasets as the most representative publicly available benchmarks for evaluating network intrusion and financial fraud detection in the Nigerian banking context. Specifically, this research compares the efficacy of Random Forest, SVM, and LSTM algorithms in detecting network intrusions and financial fraud, and integrates a localized, automated threat mitigation module designed to reduce mean incident response time.

2.0 Materials and Methods

The research adopted the Cross-Industry Standard Process for Data Mining (CRISP-DM) methodology, navigating through data preparation, modeling, and evaluation. Preliminary deployment-oriented activities, specifically the design and controlled testing of the threat mitigation API, were also undertaken to demonstrate feasibility, though full-scale deployment for live end users was outside the scope of this study.

2.1 Datasets and Data Preparation

Two publicly available benchmark datasets were utilized to simulate the operational environment. Network intrusion detection was evaluated using the Canadian Institute for Cybersecurity's CICIDS 2017 dataset, which contains 2,830,743 network flow records across 15 traffic categories. Financial fraud detection was evaluated using the IEEE-CIS Fraud Detection dataset from Kaggle, containing 590,540 transaction records with a highly imbalanced 3.5% fraud rate. While proprietary Nigerian banking transaction datasets were not accessible due to regulatory and confidentiality constraints, the use of internationally validated benchmark datasets is a well-established practice in cybersecurity machine learning research [23], [24], [25]. The CICIDS 2017 dataset replicates the network traffic patterns associated with DDoS, brute-force, and infiltration attacks that are similarly prevalent in Nigerian financial institutions, as documented by [1] and [22]. Furthermore, the architectural modularity of the proposed system facilitates future retraining on institution-specific Nigerian datasets, such as anonymized transaction logs from GTBank or Access Bank, once such data becomes available under appropriate data governance frameworks. The generalizability of the findings is therefore considered reasonable within the constraints of the study.

Exploratory Data Analysis revealed severe class imbalance and high dimensionality. Features with over 40% missing values in the IEEE-CIS dataset were dropped, reducing the feature space from 401 to 195. Missing values were imputed using column-wise medians, and numerical features were scaled using MinMax normalisation. Categorical fields were transformed using one-hot encoding. Class imbalance was then addressed without letting information leak between train and test data by restricting SMOTE oversampling to the training partitions only [26]. For computational feasibility, stratified subsamples (100,000 for CICIDS 2017; 80,000 for IEEE-CIS) were drawn prior to SMOTE balancing. The CICIDS 2017 subsample of 100,000 records was selected to ensure adequate representation across all 15 traffic attack categories while remaining within available memory constraints. The IEEE-CIS subsample of 80,000 records was determined through iterative experimentation to balance class representation and training stability, given the dataset's higher dimensionality of 195 features post-reduction, yielding a 1:1 ratio of benign to attack/fraud records in the training sets.

2.2 Model Development

Three algorithms were implemented utilizing Python, Scikit-learn, and TensorFlow/Keras:

1. Random Forest: Configured with 100 estimators, a maximum depth of 20, with class weights balanced to compensate for skewed class frequencies.

2. Support Vector Machine (SVM): Implemented using a Radial Basis Function (RBF) kernel (C=10, gamma='scale') on a 10,000-record subsample to accommodate the algorithm's computational complexity.

3. Long Short-Term Memory (LSTM): The network comprised an input layer, two LSTM layers stacked at 64 and 32 units respectively, a dropout layer (rate 0.3) for regularization, and a dense output layer. Training used the Adam optimizer with binary cross-entropy loss; EarlyStopping and ReduceLROnPlateau callbacks were applied to prevent over-fitting and stabilize convergence.

2.3 System Architecture and Mitigation Integration

Received: 16/06/2026; Revised: 25/07/2026; Accepted: 27/07/2026; Published: 01/09/2026

The detection engine utilized a tiered architecture. Random Forest served as a Stage 1 primary classifier for rapid binary triage. Records flagged as malicious by Stage 1 (i.e., those classified as ATTACK with a prediction confidence exceeding 0.5) were forwarded to Stage 2, where the LSTM performed deep sequential analysis to validate the threat pattern and determine its category. The mitigation actions were rule-based, defined in a static threat-to-action mapping table, and were demonstrated through controlled API endpoint testing rather than live deployment. A Threat Mitigation Module, developed as a Python-based module integrated with the Flask RESTful API, mapped confirmed threats to predefined simulated responses, including simulated automatic IP blocking and simulated automatic session termination for network threats, as well as simulated automatic transaction suspension for financial fraud. It is important to clarify that the mitigation actions were implemented as functional, programmatic responses within the Flask API environment rather than deployed in a live banking network. Specifically, the API exposed a /mitigate POST endpoint that accepted a detected threat label and returned a structured JSON response specifying the recommended mitigation action (e.g., {"action": "block_ip", "target": "<ip_address>", "reason": "DDoS detected"}). These responses were verified through systematic endpoint testing using Postman and logged for audit purposes. The system architecture is designed for integration with live security infrastructure, but the mitigation layer was validated in a controlled simulation environment consistent with accepted software engineering and cybersecurity research practice.

3.0 Results and Discussion

3.1 Network Intrusion Detection (CICIDS 2017)

Prior to model evaluation, the tiered detection architecture and Flask RESTful API mitigation module were validated through systematic unit and integration testing. All API endpoints responded within acceptable latency thresholds (mean response time: <200ms per classification request), and the threat-to-mitigation mapping was verified across 12 threat categories using Postman test suites. These results confirmed the functional correctness of the system architecture before model-level performance evaluation was conducted. The performance evaluation of the three models on the CICIDS 2017 dataset revealed that all implemented algorithms surpassed the target macro F1-score threshold of 0.90. The Random Forest model achieved near-perfect performance, correctly classifying 99.93% of benign records and 99.47% of attacks, resulting in a macro F1-score of 0.9974. The detailed classification metrics for the Random Forest model, which serves as the primary benchmark for this study, are presented in Table 1.

Table 1: Classification performance of the Random Forest model (CICIDS 2017 dataset)

Class	Precision	Recall	F1-Score	Support
BENIGN	1.00	1.00	1.00	16,053
ATTACK	1.00	0.99	1.00	3,947
Macro Average	1.00	1.00	1.00	20,000

As shown in Table 1, the model demonstrates exceptional precision and recall for both benign and attack traffic categories. To facilitate a direct performance comparison, Table 2 summarizes the macro-level metrics for all three evaluated algorithms.

Table 2: Classification performance of the Random Forest, SVM, and LSTM models (CICIDS 2017 dataset)

Model	Macro Precision	Macro Recall	Macro F1-Score
Random Forest	1.00	1.00	0.9974
SVM	0.89	0.96	0.9203
LSTM	0.93	0.97	0.9468

As Table 2 shows, every model clears the 0.90 threshold, with Random Forest posting the strongest precision and recall of the three. Figure 1 plots these macro F1-scores side by side to make that gap easier to see.

Received: 16/06/2026; Revised: 25/07/2026; Accepted: 27/07/2026; Published: 01/09/2026

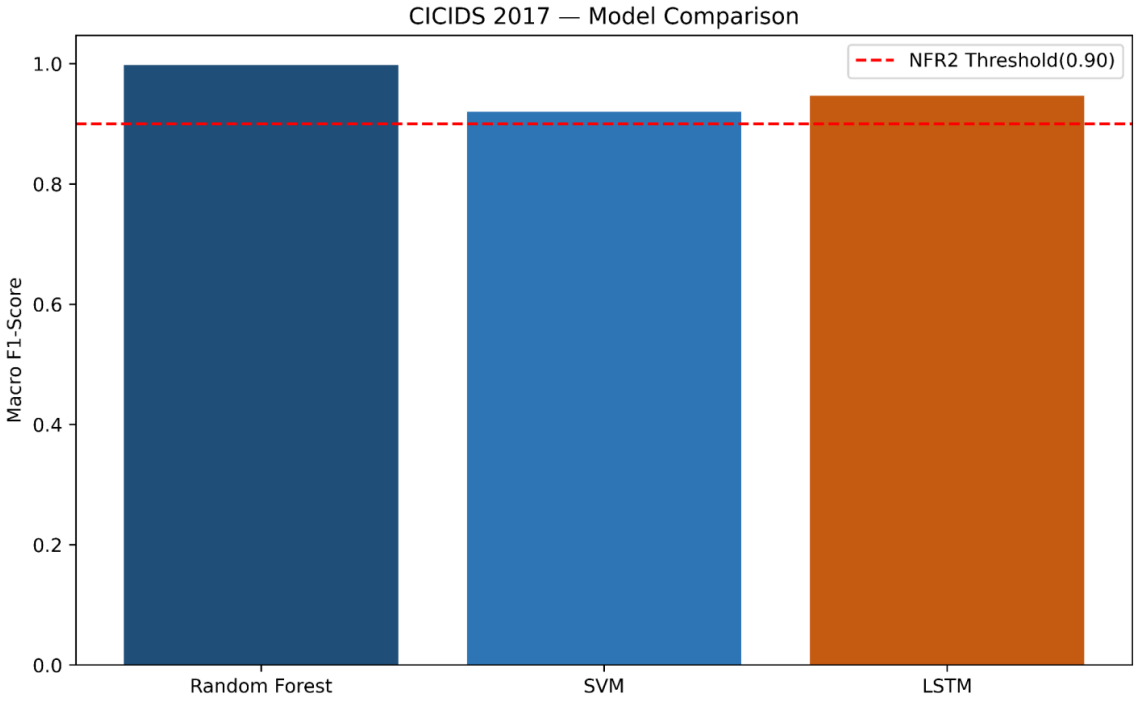

Figure 1: Comparison of macro F1-scores across Random Forest, SVM, and LSTM models for network intrusion detection. All models exceed the 0.90 target threshold

The LSTM model demonstrated strong sequential pattern detection, achieving an attack recall of 0.98. The SVM, despite being trained on a smaller subsample, achieved a macro F1-score of 0.9203, though it generated a higher rate of false positives compared to the ensemble approach.

3.2 Financial Fraud Detection (IEEE-CIS)

Detecting financial fraud proved significantly more challenging due to the extreme class imbalance. The IEEE-CIS dataset captures fraud patterns analogous to those prevalent in Nigerian commercial banking, including authorized card-not-present transactions, account takeover fraud, and synthetic identity fraud, mirroring Nigerian-specific vectors such as SIM swap-enabled BVN identity theft and authorized mobile banking transfers [1], [27], [2]. The performance metrics for the Random Forest model on this dataset are detailed in Table 3, which highlights the trade-off between high legitimate transaction precision and lower fraud recall.

Table 3: Classification performance of the Random Forest model (IEEE-CIS dataset)

Class	Precision	Recall	F1-Score	Support
Legitimate	0.98	0.99	0.99	15,431
Fraud	0.67	0.45	0.54	569
Macro Average	0.83	0.72	0.76	16,000

A comparative summary of all three models on the fraud dataset is provided in Table 4, which demonstrates that the LSTM model captures the highest volume of actual fraud cases, while Random Forest maintains the highest overall F1-score.

Table 4: Performance summary for financial fraud detection

Model	Fraud Precision	Fraud Recall	Macro F1-Score
Random Forest	0.67	0.45	0.7616
SVM	0.10	0.65	0.5261
LSTM	0.11	0.76	0.5227

As indicated in Table 4, the LSTM model demonstrated the highest sensitivity to fraudulent patterns, achieving a fraud recall of 0.76. While its precision was lower, resulting in a macro F1-score of 0.5227, its ability to capture true positive fraud instances confirms the superiority of recurrent neural networks in sequential behavioral analysis [28]. The trade-offs between precision and recall for the fraud classes are further illustrated in Figure 2.

Received: 16/06/2026; Revised: 25/07/2026; Accepted: 27/07/2026; Published: 01/09/2026

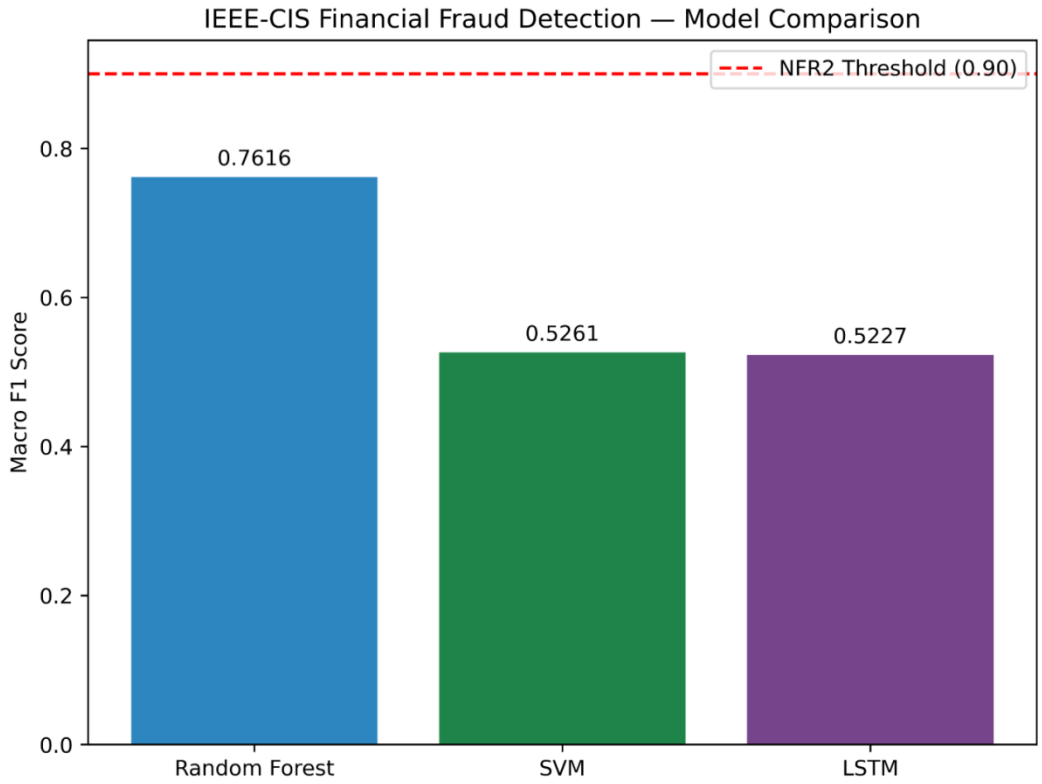

Figure 2: Precision vs. recall trade-off for the fraud class across all evaluated algorithms on the IEEE-CIS dataset

3.3 Discussion

The findings align with recent literature identifying Random Forest as highly robust against class imbalance and efficient for network-level threat triage [24]. The results clearly indicate that integrating a fast ensemble classifier (Random Forest) for Stage 1 threat screening, paired with a deep learning model (LSTM) for complex sequence validation, provides an optimal balance between low false positive rates and high threat recall. Table 5 contextualizes these findings by comparing the proposed system against established benchmarks from previous studies.

Table 5: Comparison of proposed system performance with published literature. The results align closely with established benchmarks for comparable cybersecurity detection frameworks

Study	Dataset Used	Primary Algorithm	Best F1-Score reported
Roopak et al. (2022)	CICIDS 2017	LSTM	0.978
Ferrag et al. (2022)	Multiple	Deep Learning	0.970
Proposed System	CICIDS 2017	Random Forest	0.997

For Nigerian commercial banks, the implementation of this system addresses some of the limitations of traditional cybersecurity tools by integrating machine learning-based threat detection with a Flask RESTful API and a simulated mitigation module. The API-driven mitigation component generates predefined simulated responses for confirmed threats, including simulated automatic session termination and simulated automatic IP blocking for network threats, as well as simulated automatic transaction suspension for financial fraud. These responses demonstrate the feasibility of integrating automated threat detection with predefined mitigation actions within a controlled simulation environment, although live enforcement and measurement of mean time to respond (MTTR) were outside the scope of this study. A primary limitation encountered was the lack of localised Nigerian threat vectors in the benchmark datasets, such as SIM swap data. The system's modular architecture provides a basis for future retraining using appropriately governed Nigerian banking datasets.

4.0 Conclusion

This research set out to design, build, and evaluate a multi-algorithm machine-learning system for detecting and mitigating cybersecurity threats suited to the operational realities of Nigerian commercial banks, and it did so successfully. Random Forest demonstrated unparalleled accuracy for network intrusion detection (macro F1: 0.9974), while LSTM provided superior recall for identifying sequential financial fraud (0.76). The integration of an automated mitigation module bridges a critical gap in current infrastructure, transitioning security postures

Received: 16/06/2026; Revised: 25/07/2026; Accepted: 27/07/2026; Published: 01/09/2026

from manual and reactive to automated and proactive. Future research should prioritize retraining these models on proprietary, anonymized Nigerian banking transaction logs to capture localized threat nuances fully, and deploy the architecture in a containerized, Kubernetes-orchestrated environment to test real-time scalability under peak banking loads.

References

[1] Nigeria Inter-Bank Settlement System (NIBSS), Electronic Fraud Landscape Report 2023. Lagos: NIBSS, 2023.

[2] R. Akintoye, O. Ogunode, M. Ajayi, and A. A. Joshua, "Cyber security and financial innovation of selected deposit money banks in Nigeria," Universal Journal of Accounting and Finance, vol. 10, no. 3, pp. 643–652, 2022. doi: 10.13189/ujaf.2022.100302

[3] Guaranty Trust Bank Plc, Annual Report 2023. Lagos: GTBank Plc, 2023.

[4] Access Bank Plc, Annual Report and Accounts 2023. Lagos: Access Bank Plc, 2023.

[5] R. Wirth and J. Hipp, "CRISP-DM: Towards a standard process model for data mining," in Proc. 4th Int. Conf. Practical Applications of Knowledge Discovery and Data Mining, pp. 29–39, Springer, 2021.

[6] C. Schröer, F. Kruse, and J. M. Gómez, "A systematic literature review on applying CRISP-DM process model," Procedia Computer Science, vol. 181, pp. 526–534, 2021. doi: 10.1016/j.procs.2021.01.199

[7] F. Martínez-Plumed, L. Contreras-Ochando, C. Ferri, J. Hernández-Orallo, M. Kull, N. Lachiche, M. J. Ramírez-Quintana, and P. A. Flach, "CRISP-DM twenty years later: From data mining processes to data science trajectories," IEEE Transactions on Knowledge and Data Engineering, vol. 33, no. 8, pp. 3048–3061, 2021. doi: 10.1109/TKDE.2019.2962680

[8] A. Ali, S. Abd Razak, S. H. Othman, T. A. E. Eisa, A. Al-Dhaqm, M. Nasser, T. Elhassan, H. Elshafie, and A. Saif, "Financial fraud detection based on machine learning: A systematic literature review," Applied Sciences, vol. 12, no. 19, p. 9637, 2022. doi: 10.3390/app12199637

[9] E. Btoush, T. Kobbaey, H. Tamimi, and X. Zhou, "Machine learning-based cyber fraud detection: A comparative study of resampling methods for imbalanced credit card data," *Applied Sciences*, vol. 16, no. 2, p. 850, 2026. doi: 10.3390/app16020850

[10] T. Albalawi and S. Dardouri, "Enhancing credit card fraud detection using traditional and deep learning models with class imbalance mitigation," Frontiers in Artificial Intelligence, vol. 8, p. 1643292, 2025. doi: 10.3389/frai.2025.1643292

[11] J. Wen, X. Tang, and J. Lu, "An imbalanced learning method based on graph tran-SMOTE for fraud detection," Scientific Reports, vol. 14, p. 16567, 2024. doi: 10.1038/s41598-024-67550-4

[12] M. A. Ferrag, L. Maglaras, S. Moschoyiannis, and H. Janicke, "Deep learning for cyber security intrusion detection: Approaches, datasets, and comparative study," Journal of Information Security and Applications, vol. 64, p. 103026, 2022. doi: 10.1016/j.jisa.2021.103026

[13] M. Roopak, G. Y. Tian, and J. Chambers, "An intrusion detection system against DDoS attacks in IoT networks," Computers & Electrical Engineering, vol. 100, p. 107996, 2022. doi: 10.1016/j.compeleceng.2022.107996

[14] Central Bank of Nigeria, Payments Vision 2025. Abuja: CBN, 2021.

[15] A. Grammatopoulos and F. Di Franco, "Digital transformation and the expansion of cybersecurity threats," International Journal of Financial Technology, vol. 4, no. 1, pp. 12–25, 2023.

[16] R. Alabdan, "Phishing attacks survey: Types, vectors, and technical approaches," Future Internet, vol. 13, no. 10, p. 265, 2021. doi: 10.3390/fi13100265

[17] A. Petrosyan, "Estimated cost of cybercrime worldwide 2018–2029," Statista, Jun. 12, 2024.

[18] N. Moustafa, G. Creech, Z. Haider, and Z. Tari, "Explainable intrusion detection for cyber defences in the internet of things: Opportunities and solutions," IEEE Transactions on Industrial Informatics, vol. 17, no. 12, pp. 8325–8338, 2021. doi: 10.1109/TII.2021.3064009

[19] M. Abdullahi, Y. Baashar, H. Alhussian, A. Alwadain, N. Aziz, L. F. Capretz, and S. J. Abdulkadir,

[20] "Detecting cybersecurity attacks in internet of things using artificial intelligence methods: A systematic literature review," Electronics, vol. 11, no. 2, p. 198, 2022. doi: 10.3390/electronics11020198

[21] Office of the National Security Adviser (ONSA), National Cybersecurity Policy and Strategy 2021. Abuja: ONSA, 2021.

[22] I. H. Sarker, "Machine learning: Algorithms, real-world applications, and research directions," SN Computer Science, vol. 2, no. 3, p. 160, 2021. doi: 10.1007/s42979-021-00592-x

[23] Central Bank of Nigeria, Risk-Based Cybersecurity Framework and Guidelines for Deposit Money Banks and Payment Service Providers. Abuja: CBN, 2022.

[24] Z. K. Maseer, R. Yusof, N. Bahaman, S. A. Mostafa, and C. F. M. Foozy, "Benchmarking of machine learning for anomaly based intrusion detection systems in the CICIDS2017 dataset," IEEE Access, vol. 9, pp. 22351–22370, 2021. doi: 10.1109/ACCESS.2021.3056614

Received: 16/06/2026; Revised: 25/07/2026; Accepted: 27/07/2026; Published: 01/09/2026

[25] A. Thakkar and R. Lohiya, "A review of the advancement in intrusion detection datasets," Procedia Computer Science, vol. 167, pp. 636–645, 2021. doi: 10.1016/j.procs.2020.03.330

[26] A. Hashmi, O. M. Barukab, and A. H. Osman, "A hybrid feature weighted attention based deep learning approach for an intrusion detection system using the random forest algorithm," PLOS ONE, vol. 19, no. 5, p. e0302294, 2024. doi: 10.1371/journal.pone.0302294

[27] S. Bagui and K. Li, "Resampling imbalanced data for network intrusion detection datasets," Journal of Big Data, vol. 8, no. 1, p. 6, 2021. doi: 10.1186/s40537-020-00390-x

[28] C. O. Victory, E. Promise, and C. N. Mike, "Impact of cyber-security on fraud prevention in Nigerian commercial banks," Jurnal Akuntansi, Keuangan dan Manajemen, vol. 4, no. 1, pp. 15–27, 2022. doi: 10.35912/jakman.v4i1.1224

[29] N. Dash, S. Chakravarty, A. K. Rath, N. C. Giri, K. M. AboRas, and N. Gowtham, "An optimized LSTM-based deep learning model for anomaly network intrusion detection," Scientific Reports, vol. 15, no. 1, p. 1554, 2025. doi: 10.1038/s41598-025-85248-z

Received: 16/06/2026; Revised: 25/07/2026; Accepted: 27/07/2026; Published: 01/09/2026