Case Study
ISSN (Online): 3048-894X

Journal of Digital Security and Forensics
July-December 2025 2(2), 137-143

LEVERAGING HYBRID AI FOR REAL-TIME FRAUD DETECTION: A CASE STUDY ON THE EFFICACY OF GRAPH NEURAL NETWORKS AND ANOMALY DETECTION IN NIGERIAN FINTECHS

Temitayo Oluwaseun Jejeniwa [1], Titilola Olaide Jejeniwa [2], Olugbenga Saheed Owolabi [3]

[1] United Nations African Centre for Space Science Technology Education-English, NASRDA, Obafemi Awolowo University, Ile Ife, Osun State, Nigeria
[2] Advanced Space Technology Laboratory (Southwest), National Space Research and Development Agency, Obafemi Awolowo University, Ile Ife, Osun State, Nigeria
[3] Itap Solutions Limited, Suite 23, Block A, Alausa Shopping Mall, Ikeja Lagos, Nigeria

Received 15 October 2025
Accepted 13 November 2025
Published 19 December 2025

Corresponding Author

Temitayo Oluwaseun Jejeniwa,
jejeniwatemitayo@yahoo.com

DOI
10.29121/DigiSecForensics.v2.i2.2025.60

Funding: This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

ABSTRACT

This study investigates the development and application of a hybrid artificial intelligence (AI) model for real-time fraud prevention within Nigeria's rapidly growing FinTech sector. The research addresses the critical challenge of sophisticated financial fraud, which hampers financial inclusion and erodes consumer trust. Moving beyond traditional single-model approaches, this paper proposes a novel framework integrating Graph Neural Networks (GNNs) to analyze complex transactional relationships and an Isolation Forest algorithm for point anomaly detection. Using a real-world, anonymized transaction dataset from a major Nigerian FinTech company, the study trains and validates the hybrid model. Key performance metrics (Precision, Recall, F1-Score, AUC-ROC) are evaluated against benchmark models, including Logistic Regression and Random Forest. The results demonstrate the superior efficacy of the hybrid AI approach, achieving an F1-score of 0.92 and an AUC-ROC of 0.98, significantly outperforming the benchmarks in accurately flagging fraudulent transactions while minimizing false positives. The study concludes that a hybrid model is particularly suited for the unique challenges of the Nigerian FinTech landscape and provides strategic recommendations for the practical integration of such explainable AI systems to bolster security and foster sustainable growth.

Keywords: Financial Fraud, Fintech, Artificial Intelligence, Graph Neural Networks, Anomaly Detection, Real-Time Systems, Nigeria, Hybrid Model

1. INTRODUCTION

The Nigerian financial technology (FinTech) sector has emerged as a pivotal force in driving financial inclusion, leveraging mobile technology to provide accessible banking, payment, and lending services to a previously underserved

population Central Bank of Nigeria. (2022). However, this rapid digitization and growth have been paralleled by a surge in sophisticated financial fraud, including identity theft, account takeover, and transaction laundering. The Economic and Financial Crimes Commission Economic and Financial Crimes Commission (2023) reported a significant increase in digitally-enabled financial crimes, resulting in billions of Naira in losses annually and posing a substantial threat to consumer trust and sector stability.

Traditional rule-based fraud detection systems and conventional machine learning models like Logistic Regression often fall short in this evolving landscape. They struggle to detect coordinated, networked fraud (e.g., sleeper cells, money muling networks) and are frequently plagued by high false positive rates, which degrade customer experience Jurgovsky et al. (2018). Furthermore, a significant challenge in the Nigerian context is the scarcity of large, high-quality, labeled datasets for training supervised models, as fraud patterns are dynamic and often go unreported or unlabeled.

This study addresses these gaps by proposing and evaluating a novel hybrid AI model tailored for real-time fraud detection in Nigerian FinTechs. The novelty of this research is fourfold: (1) it combines the relational analysis power of Graph Neural Networks (GNNs) with the unsupervised strength of an Isolation Forest algorithm to detect both known and novel fraud patterns; (2) it focuses on real-time implementation, a critical requirement for the high-velocity Nigerian market; (3) it is specifically designed for the unique data and operational environment of Nigerian FinTechs; and (4) it offers a practical solution to the problem of data scarcity by leveraging a hybrid supervised-unsupervised approach.

The primary research objectives are:

1) To design a hybrid AI fraud detection model architecture combining GNNs and an Anomaly Detection model.
2) To procure, preprocess, and structure a real-world transactional dataset from a Nigerian FinTech.
3) To implement, train, and evaluate the proposed hybrid model against established benchmark models.
4) To compare model performance based on accuracy, precision, recall, F1-score, and false positive rates.
5) To propose a framework for the practical, real-time integration of the hybrid model into FinTech financial management systems.

2. LITERATURE REVIEW

Financial fraud detection has evolved from manual reviews and simple heuristic rules to sophisticated machine learning (ML) and deep learning (DL) models. Early academic and industrial work heavily featured models like Logistic Regression, Decision Trees, and ensemble methods like Random Forest and XGBoost Bahnsen et al. (2016). These models primarily treat transactions as independent and identically distributed (i.i.d.) events, focusing on features intrinsic to a single transaction (e.g., amount, time, merchant category).

However, fraud is inherently a relational problem. Criminals often operate in networks, sharing resources like devices, bank accounts, and beneficiary details. This limitation has spurred interest in graph-based analytics. Techniques such as community detection and label propagation on graphs have been used to identify

suspicious clusters Pandey et al. (2021). The recent advent of Graph Neural Networks (GNNs) represents a significant leap forward. GNNs can learn powerful representations of nodes (e.g., users, devices) and edges (transactions) within a graph, capturing complex relational patterns that other models miss Liu et al. (2021). Their application in finance, however, remains nascent, especially in emerging markets.

Simultaneously, the need to detect never-before-seen fraud patterns has led to the adoption of unsupervised anomaly detection algorithms like Isolation Forest Liu et al. (2008) and Autoencoders. These models learn a profile of "normal" behavior and flag significant deviations, making them ideal for identifying novel attacks without relying on labeled fraud data.

A few studies have begun exploring hybrid models. For instance, Roy et al. (2022) combined a GNN with a Gradient Boosting Machine, but their focus was on offline analysis rather than real-time performance. The specific application of a GNN-Anomaly detection hybrid, particularly within the high-stakes, data-scarce, and rapidly evolving context of an African FinTech ecosystem like Nigeria's, constitutes a clear gap in the literature. This study aims to fill this gap by building and validating a integrated framework that leverages the strengths of both approaches for real-time defense.

3. METHODOLOGY

3.1. RESEARCH DESIGN

This study adopted a quantitative, design-science approach. The core activity was the construction (design) and rigorous evaluation of a novel AI artifact—the hybrid fraud detection model—in response to a identified and relevant business problem Hevner et al. (2004).

3.2. DATA COLLECTION AND DESCRIPTION

Through a research partnership with a leading Nigerian FinTech company (hereafter referred to as "FinTech-NG" for anonymity), access was granted to an anonymized dataset of 1.5 million transactions over a six-month period. The dataset included:

- **Transaction attributes:** transaction ID, hashed user ID, timestamp, amount, location (state), merchant category code (MCC), device ID (hashed), IP address (anonymized), beneficiary account number (hashed), transaction status.
- **Label:** A binary flag indicating confirmed fraudulent (1) or legitimate (0) transactions, as determined by FinTech-NG's internal fraud investigation team and chargeback records. The dataset was highly imbalanced, with a fraud rate of 0.8%.

3.3. DATA PREPROCESSING

The data underwent standard preprocessing:

1) **Handling Missing Values:** Categorical missing data (e.g., MCC) was imputed with a new "unknown" category. Numerical missing data (e.g., age) was imputed with the median.
2) **Feature Engineering:** New temporal features were created, including transaction hour, day of the week, and a rolling window feature for a

user's transaction frequency (number of transactions in the last 24 hours).

3) **Normalization:** Numerical features (amount, transaction frequency) were standardized using Z-score normalization.
4) **Graph Construction:** For the GNN component, a heterogeneous graph was constructed with two node types: User and Device. Edges represented transacted_with relationships between users and devices. Node features for User included age bracket and account age. Edge features included transaction amount and time-derived features.

3.4. MODEL DEVELOPMENT

The hybrid model consisted of two parallel components whose scores were fused for a final decision.

1) **Component A:** Graph Neural Network (Relational Model). A GraphSAGE model was implemented to learn embeddings for each user and device node. The model learned to aggregate features from a node's neighbors to generate a representation that encapsulates local graph structure. This representation was passed through a multilayer perceptron (MLP) to produce a probability score (score_gnn) indicating the likelihood of a transaction being fraudulent based on its relational context.
2) **Component B:** Isolation Forest (Anomaly Model). An Isolation Forest model was trained on the preprocessed transactional feature set (amount, frequency, location, etc.) for all users. This model learned to isolate point anomalies and output an anomaly score (score_anom) for each transaction.
3) **Hybridization:** The scores from both components (score_gnn, score_anom) for each transaction were combined into a final feature vector. A meta-classifier—a Logistic Regression model—was then trained on this two-dimensional feature vector, with the original labels as targets, to learn the optimal way to combine the two scores into a single, final fraud probability.

3.5. BENCHMARK MODELS

The hybrid model's performance was compared against two strong benchmark models:

1) **Random Forest:** A powerful ensemble method known for robust performance on tabular data.
2) **XGBoost:** A highly efficient and effective gradient boosting framework, often a top performer in machine learning competitions.

3.6. EVALUATION METRICS

Given the severe class imbalance, accuracy was deemed a misleading metric. Models were evaluated primarily on:

- **Precision:** The proportion of predicted frauds that were actual frauds.
- **Recall (Sensitivity):** The proportion of actual frauds that were correctly identified.

- **F1-Score:** The harmonic means of Precision and Recall.
- **Area Under the Receiver Operating Characteristic Curve (AUC-ROC):** The model's ability to distinguish between classes across all thresholds.
- **False Positive Rate (FPR):** The proportion of legitimate transactions incorrectly flagged as fraud.

4. RESULTS

The dataset was split into 70% training, 15% validation, and 15% testing. The models were trained on the training set, hyperparameters were tuned on the validation set, and final performance was reported on the held-out test set.

Table 1

Table 1 Performance Comparison of Fraud Detection Models on the Test Set

Model	Precision	Recall	F1-Score	AUC-ROC	False Positive Rate
Random Forest	0.85	0.8	0.82	0.93	0.015
XGBoost	0.87	0.83	0.85	0.95	0.013
Proposed Hybrid	**0.94**	**0.9**	**0.92**	**0.98**	**0.008**

As shown in Table 1, the proposed hybrid model significantly outperformed both benchmark models across all key metrics. It achieved the highest F1-score (0.92) and AUC-ROC (0.98), indicating a superior overall balance between identifying true frauds and avoiding false alarms. Crucially, it achieved the lowest False Positive Rate (0.008), which is a critical operational metric for minimizing customer disruption.

Analysis of the GNN component revealed it was particularly effective at identifying clustered fraud, such as multiple user accounts linked to a single device making rapid, small transactions. The Isolation Forest component excelled at flagging dramatic behavioral shifts, such as a typically low-activity user suddenly initiating a very large transfer.

5. DISCUSSION

The results strongly support the thesis that a hybrid AI approach is superior for real-time fraud detection in the Nigerian FinTech context. The synergy between the GNN and the anomaly detector proved powerful: the GNN exposes fraud based on who you are connected to, while the Isolation Forest exposes fraud based on how you are behaving differently. This dual lens is essential for combating the multifaceted nature of financial fraud.

The low false positive rate of the hybrid model is a major practical advantage. For a FinTech, a high false positive rate directly translates to blocked legitimate transactions, customer support costs, and ultimately, customer churn. By minimizing this, the model enhances both security and user experience.

This study also demonstrates a viable path forward in environments with labeled data scarcity. The unsupervised anomaly detection component continuously learns "normal" behavior and can flag novel fraud patterns that have not yet been labeled, making the system more adaptive and resilient to new attack vectors.

6. IMPLICATIONS FOR PRACTICE

For Nigerian FinTechs, the implementation framework involves:

1) **Data Pipeline:** Establishing a real-time stream of transaction data to a graph database and a feature store.
2) **Model Serving:** Deploying the trained GNN and Isolation Forest models as microservices that can score transactions in milliseconds.
3) **Decision Engine:** Implementing the meta-classifier to combine the scores and make a final decision based on a tunable threshold aligned with the company's risk appetite.
4) **Feedback Loop:** Creating a mechanism to feed confirmed fraud labels (from investigations) back into the system to continuously retrain and improve the models.

7. LIMITATIONS AND FUTURE RESEARCH

This study has limitations. The data came from a single FinTech, and while representative, generalizability to all Nigerian financial institutions could be further validated. Furthermore, the graph constructed was relatively simple; future work could explore more complex graphs incorporating merchant nodes, IP addresses, and geographic locations.

Future research directions include: (1) incorporating explainable AI (XAI) techniques to provide clear reasons for a flagged transaction, aiding investigators; (2) exploring online learning techniques to allow the model to adapt continuously without full retraining; and (3) extending the framework to other emerging markets with similar characteristics.

8. CONCLUSION

This study successfully designed, implemented, and validated a hybrid AI model for real-time fraud detection tailored to the Nigerian FinTech sector. By integrating the relational intelligence of Graph Neural Networks with the deviation-detection capabilities of an Isolation Forest algorithm, the model demonstrated a significant performance improvement over traditional, single-model approach. It achieved higher fraud detection accuracy while critically maintaining a lower false positive rate.

The findings provide a compelling case for Nigerian FinTechs to invest in advanced, hybrid AI systems. Such systems are not merely technological upgrades but are foundational to building a secure, trustworthy, and inclusive digital financial ecosystem. As fraudsters become more sophisticated, the defense must evolve beyond siloed solutions towards integrated, intelligent, and real-time frameworks like the one proposed here.

CONFLICT OF INTERESTS

None .

ACKNOWLEDGMENTS

None.

REFERENCES

Bahnsen, A. C., Aouada, D., and Ottersten, B. (2016). Feature Engineering Strategies for Credit Card Fraud Detection. Expert Systems with Applications, 51, 134–142. https://doi.org/10.1016/j.eswa.2015.12.030

Central Bank of Nigeria. (2022). Payments System Vision 2025.

Economic and Financial Crimes Commission. (2023). EFCC Annual Report 2022.

Hevner, A. R., March, S. T., Park, J., and Ram, S. (2004). Design Science in Information Systems Research. MIS Quarterly, 28(1), 75–105. https://doi.org/10.2307/25148625

Jurgovsky, J., Granitzer, M., Ziegler, K., Calabretto, S., Portier, P. E., He-Guelton, L., and Caelen, O. (2018). Sequence Classification for Credit-Card Fraud Detection. Expert Systems with Applications, 100, 234–245. https://doi.org/10.1016/j.eswa.2018.01.037

Liu, F. T., Ting, K. M., and Zhou, Z. H. (2008). Isolation Forest. In Proceedings of the 2008 Eighth IEEE International Conference on Data Mining (pp. 413–422). IEEE. https://doi.org/10.1109/ICDM.2008.17

Liu, Y., Li, Z., Zhou, C., Jiang, Y., Sun, J., Wang, M., and He, X. (2021). Generative Adversarial Networks for Anomaly Detection in Financial Time Series. ACM Transactions on Knowledge Discovery from Data, 15(4), Article 60, 1–23.

Pandey, A., Siripurapu, A., and Kumar, S. (2021). A Graph-Based Approach for Financial Fraud Detection. In 2021 IEEE International Conference on Big Data (Big Data) (pp. 1759–1768). IEEE.

Roy, A., Sun, J., Mahoney, R., Alon, L., Jin, M., and Fletcher, T. (2022). Deep Learning for Fraud Detection in Financial Graphs. ACM Computing Surveys, 55(4), Article 79, 1–26. https://doi.org/10.1145/3468266