Advances in Knowledge-Based Systems, Data Science, and Cybersecurity; Research Received 04-02-2026; Accepted 06-04-2026; Published online 30-04-2026

Anomaly Detection Frameworks for Cybersecurity in Nigerian Fintech Startups Using Unsupervised Machine Learning Techniques

Omowunmi O Adebajo ooadebajo@bellsuniversity.edu.ng
Department of Computer Engineering
Bells University of Technology.
Ota, Ogun State, Nigeria

Surajudeen A Babatunde babatundesurajudeen@funaab.edu.ng
Department of Mechatronics Engineering
Federal University of Agriculture.
Abeokuta, Ogun State, Nigeria

Akeem A Oni oniaa@bellsuniversity.edu.ng
Department of Computer Engineering
Bells University of Technology.
Ota, Ogun State, Nigeria

Johnson O Abiola joabiola@bellsuniversity.edu.ng
Mechatronics Engineering Department,
Bells University of Technology.
Ota, Ogun State, Nigeria

Oluwaseye A Adebajo oluwaseye.adebajo@bowen.edu.ng
Anatomy Programme,
College of Health Sciences
Bowen University Iwo Campus
Osun State, Nigeria

Corresponding Author: Omowunmi Olabowale Adebajo

Abstract

This research establishes an unsupervised machine learning-based anomaly detector that will enhance cybersecurity among Nigerian fintech startups where high rates of digital growth, volumes of transactions, and critical shortages of labeled fraud data render traditional rule-based and supervised fraud detection systems ineffective. The need to have adaptive and scalable security systems capable of identifying emerging and novel fraud patterns and keeping computational efficiency and operational feasibility in resource-constrained fintech contexts drives the study. The study employs the real-life experience of actual transactions with a financial technology application in Nigeria to model standard transactional behavior and identify abnormalities as a deviation of the standard. It does it in a systematic process, which includes data preprocessing, feature normalization and behavioral analysis. Three unsupervised models were used, namely, clustering-based approaches (K-Means), isolation-based

Citation: Adebajo, et al. Anomaly Detection Frameworks for Cybersecurity. **Advances in Knowledge-Based Systems, Data Science, and Cybersecurity**. 2026;3(1).23.

Page number:478-49

approaches (Isolation Forest), and reconstruction-based approaches (PCA-based reconstruction). They have been selected, as they are complementary theories that are also applicable to the Nigerian fintech situation. They were compared according to ROC-AUC, the values of anomaly scores, the false-positive results, and their performance. The results show that models based on isolation and reconstruction always do better than simple clustering methods. The ROC-AUC of Isolation Forest on the training set and held-out test set was 0.60 and 0.57, respectively, indicating that the generalization is relatively stable over the other models. Models based on reconstruction were more stable and less prone to false-positive rates, however. Overall, the results show that unsupervised anomaly detection is a useful and effective way to improve cybersecurity in fintech. It enables the early identification of the suspicious activity without requiring labelled information and offers a scalable framework to support fraud prevention, regulatory compliance, and secure online financial transactions within the Nigerian fintech ecosystem.

Keywords: Cybersecurity, Anomaly detection, Fintech, Nigeria, Unsupervised machine learning, Isolation Forest, Autoencoder, Fraud detection.

1. INTRODUCTION

The fintech sector has rapidly expanded globally in recent years, with even more accelerated growth in Nigeria. This has facilitated access to financial services via digital platforms [1, 2] . The increase in digital transactions facilitates the occurrence of fraud, money laundering, and many cybersecurity issues. Considering the high rate at which the volume of transactions is increasing, traditional rule-based detection solutions find it difficult to cope with the new risks. Hence, there is a need for scalable, automated, and adaptive ways to find anomalies [3].

Machine learning has emerged as a highly useful method of fraud detection in the financial system due to its ability to search through large volumes of transaction data and single out an unusual pattern (discontinuous), which could be deemed as fraudulent [3]. Systematic review provided by [4], shows that machine-learning techniques, in which a supervisor participates, are actively used to detect fraud; however, the lack of labelled instances of fraud is a serious problem since the number of fraud transactions is minimal compared to the general activity. This has in turn seen the unsupervised learning methods take the center stage. Clustering algorithms, autoencoders, and isolation-based algorithms can be trained to discover the norm behavior of transactions using large amounts of unlabeled data and then indicate any deviations as possible anomalies [5]. Recent research has introduced ensemble unsupervised methodologies that integrate autoencoders, Self-Organizing Maps (SOMs), and Restricted Boltzmann Machines (RBMs), showcasing superior accuracy and a marked decrease in false positives relative to conventional unsupervised benchmarks [6]. Financial transactions often create multivariate and temporal sequences with important dependencies across features and time, in addition to static transaction snapshot analyses. Recent research on multivariate timeseries anomaly detection has put forward advanced architectures that integrate neural networks with temporal modeling, facilitating the identification of subtle yet significant deviations that may develop over time [7]. Even though these methods have gotten better, it's still hard to put them into practice, especially in Nigerian fintech startups that don't have a lot of money. A few of the key issues include lack of sufficient engineers and security personnel, lack of sufficient computing resources and ability to notice things in real time with low latency. Also, rules about anti-money

laundering (AML) require traceability, auditability, and minimizing false alarms to keep legitimate business operations from being disrupted.

In Nigeria, there is growing evidence that unsupervised learning can effectively identify credit card and transaction fraud, even in the absence of substantial labeled fraud data. This has been supported by [8] and [9], which show that AI-based anomaly detection models on banking institutions in Nigeria can discover fraudulent behavior without using pre-labeled corpora; by [10], which also reports that unsupervised pipelines trained on Nigerian fintech data achieve significant discrimination; and However, several current studies still rely on benchmark or global data rather than locally obtained fintech transaction data [11]. This gap underscores the necessity of a research that is expressly designed to respond to the reality of Nigerian fintech activities. The given gap underlines the relevance of the research which is obviously sensitive to the reality of the fintech work in Nigeria. Based on that, the current research aims to fill this gap with the help of actual transaction statistics of a Nigerian fintech company outlining normative transaction practices, deploying a range of unsupervised anomaly detection algorithms, and analyzing their efficiency. The ultimate goal is to implement an acceptable anomaly-detection service to the financial companies in Nigeria, which can meet an acceptable trade-off among accuracy, false-positive rates, computational efficiency, interpretability and compliance standards.

The aim of the research is to investigate and rigorously evaluate a machine-learning algorithm of anomaly detection without supervision, thereby enhancing cybersecurity resilience of Nigerian fintech firms.

This research is based on the concept and testing of an anomaly detection system that is based on unsupervised learning algorithms, and the aim is to supplement cybersecurity measures in the rapidly expanding industry of Nigerian financial technology companies. The research is based on three overarching research objectives: (i) to perform an intensive study of the available fintech transaction samples to establish baseline behavioural heuristics and salient statistical characteristics that support the concept of anomaly detection; (ii) to architect and operationalize a set of unsupervised learning models or algorithms such as (a) clustering, (b) isolation, and (c) reconstruction algorithms exploiting the curated transaction corpus; and (iii) to systematically evaluate and compare the performance of these realizations using quantitative metrics of anomaly scores, detection accuracy, false positives.

2. LITERATURE REVIEW

2.1 Cybersecurity in Digital Financial Systems

In the digital financial systems, cybersecurity refers to the set of strategies, tools, and protective systems that are applied to safeguard financial platforms, user information, transactions, and organizational resources against intrusion by hackers and data breaches, fraud, and a spectrum of other cyber threats. The field has become one of the fundamental forces of confidence, reliability, and functional integrity, especially as the services of fintech become a more and more central part of financial inclusion and daily economic practice [12]. The lure of online financial system(s) to cyber criminals is the concentration of sensitive personally identifiable information (PII), payment credentials, and highfrequency transactional information that fintech platforms usually contain.

Common attacks are phishing attacks, malware, identity theft, data breaches, insider attacks, API attacks, and advanced social engineering attacks [13].

The fintech industry in Nigeria has been growing at a very high pace over the last ten years due to regulatory changes and infrastructure development that have increased the uptake of mobile banking and online payment platforms. Enterprises having a high growth today make hundreds of millions of transactions every month, which indicates how important the involvement of Nigerian fintech providers in the state economy and the financial system as a whole is [14]. Cybersecurity issues facing the Nigerian fintech companies are not only based on the global trends of the threats but also on the national specific issues related to the system. Practitioners and scholars have reported insecure mobile attacks such as phishing, SIM-swap attacks, account takeover attacks, mobile malware, and exploitation of poorly configured APIs [15]. Numerous startups are limited in terms of resources and capabilities: the lack of qualified security staff, insufficient funds, and constant changes in regulatory demands. These limitations make these firms more susceptible and slow down the reaction to

2.2 Anomaly Detection in Cybersecurity

Anomaly detection is a methodological framework that is used to detect observations, events, or data points in a dataset that significantly deviate in relation to already existing normative patterns. Within the context of cybersecurity, these aberrations are monitored due to their possible indicators of latent error, misconfiguration, malfunctions of a system, or undesirable actions, such as fraud and hacking [15]. From literature, it routinely identifies three main types of anomalies: point anomalies are individual cases that do not follow the pattern; contextual anomalies are individual instances that do not deviate but instead follow a pattern in one environment and do not in another, like an atypical transaction time or machine use; and collective anomalies are cases where a group of data points, all of which by themselves are normal, do not conform to a pattern, such as coordinated fraudulent schemes [2, 3].

Anomaly detection is found to be particularly beneficial in financial systems, where the use of this technique is used to continually track transaction streams, user actions, account access records, and flows of payments. Abnormalities in transfers, unusual order of minor transactions with withdrawals, abrupt increases in activity, or coordinated transfers between accounts are all characteristics that can be tagged as abnormalities that require additional investigation or automated preventive measures. With a systematic detection and resolution of these abnormalities, institutions will be able to prevent fraud, money laundering, account hijacking, and a myriad of other illegal practices [16].

2.3 Machine Learning Techniques for Anomaly Detection

2.3.1 Clustering-based techniques

Clustering methods divide similar observations into discrete clusters and identify anomalies which do not fit into the assumed cluster shapes. The K-Means is computationally fast and performs well with large volumes of financial data, but it assumes that clusters are hyper-spherical and have equal variance, which is not always true, and when transaction behavior deviates from it, detecting

aberrations can become significantly worse [12]. DBSCAN addresses these issues by clustering points according to their density, identifying non-circular clusters, and designating low-density points as anomalies [17]. Dendrograms facilitate the comprehension of hierarchical clustering; yet, they are computationally intensive for extensive datasets [18].

2.3.2 Isolation-based techniques

Isolation-based methods identify outliers through recursive splitting, effectively distinguishing them from the remainder of the data. This is effective because outliers are infrequent and possess attribute values that significantly deviate from the majority of data points. Isolation Forest (iForest) randomly selects features and divide values. Anomalies require fewer divisions to achieve isolation. Seclusion Forest is effective with

high-dimensional data as it does not require the estimation of distance or density. Research regularly confirms its effectiveness in fraud detection, intrusion detection, and payment anomaly monitoring [4]. The Extended Isolation Forest (EIF) enhances the original model by utilizing random hyper-planes rather than axis-aligned splits, hence improving isolation efficacy for intricate or nonlinearly structured data [19].

2.3.3 Reconstruction-based techniques

Neural network models that are based on reconstruction learn how to make compressed versions of normal behavior and then use those to find anomalies, which are cases that can't be accurately reconstructed. Autoencoders (AEs) are a kind of neural network that has an encoder that takes data and compresses it into a hidden form, and a decoder that puts the original data back together. The network is trained to minimize the reconstruction loss to learn the basic structure of normal transaction data. The model tends to make numerous errors to reconstruct when the model observes unusual behavior [20]. Variational Autoencoders (VAEs) improve this architecture by learning a probabilistic latent distribution instead of a single deterministic embedding. This is particularly powerful in detecting anomalies, as the differences between the trained latent distribution manifest themselves as high reconstruction error [16]. In the current analysis, PCA-based reconstruction is used as the reconstruction-based model, which can capture the main variance structure of the normal transactions and identify the observations with high mean squared reconstruction error as anomalies.

2.4 Comparative Analysis of Unsupervised Techniques

Deep reconstruction-based models (e.g., autoencoders and variational autoencoders (VAEs)) tend to detect financial anomalies more accurately in empirical assessment than traditional clustering models [18]. Methods based on isolation are great for finding things and can be used on a large scale. An example is that Isolation Forest is almost linear time. Local Outlier Factor (LOF) and other density-based methods work well when anomalies are small changes in the data, but they don't work as well when there are more dimensions [12]. In the case of fintech transaction data, isolation-based methods are optimal when the data has large dimensions but few fraud cases, reconstruction-based

neural models are optimal when the data behavior is complex and occurs within a short time, and clustering-based methods fail when there are mixed/overlapping distributions of transactions.

3. METHODOLOGY

3.1 Research Design

This research is a quantitative, experimental, and data-driven one that will come up with and test anomaly detection frameworks tailored to cybersecurity usage in Nigerian fintech startups, utilizing unsupervised machine-learning methods. Following the research paradigm, which supposes that statistical processes and algorithmic methods can objectively depict abnormal or deceitful conduct, the research is conducted within a strictly structured and sequentially scheduled workflow, which comprises: (i) exploratory data analysis (EDA) to comprehend the structure of the data, attribute distribution, and data-quality issues; (ii) a data-processing pipeline; and (iii) feature FIGURE 1, shows the process of the research.

3.2 Dataset Description

The data utilized in this study is anonymized transactional data of a Nigerian fintech firm within the digital payments ecosystem. The complete data has around 50,000 transactions that occurred between 2022 and 2024 with 10,000 transactions (considered a held-out test set). Records are encoded with pseudonymized user identifiers, Unix time stamps, transaction value (in Nigerian Naira), merchant category codes, payment channels (mobile app, USSD, or web), transaction type (fund transfer, merchant payment, airtime purchase, withdrawal, deposit or inter-account micro-transaction), device identifiers, and binary fraud label based on post-hoc operational review by the fraud operations team at the fintech. The label of fraud was not employed in the training of the models, but it was provided as ground truth only during evaluation. These records have a granularity that allows the individual user-level and system-wide modeling of normative behavioral patterns.

A first glance shows general data quality issues that are characteristic of real-world financial transaction data: missing data, mismatched categorical names, redundancy, stochastic noise, and numerical outliers. Additionally, the dataset presents itself with an overwhelming class imbalance problem where fraud cases account for only 0.9% of total transactions while the remaining 99.1% are normal transaction cases. This large imbalance between classes is enough to warrant unsupervised learning methods be applied to detect anomalies with the help of Isolation Forest, K-Means clustering, and PCA reconstruction, which do not always need to be fed with labelled cases of fraud to learn. Personally identifiable information had been stripped off the data before the start of analysis; user ids had been assigned pseudonymously with hash tokenization to meet the NDPR 2023 guidelines. Time richness of this data assists in extracting higher-order features, such as frequency of transactions, time-of-day variables, repetitive weeks and gaps between transactions.

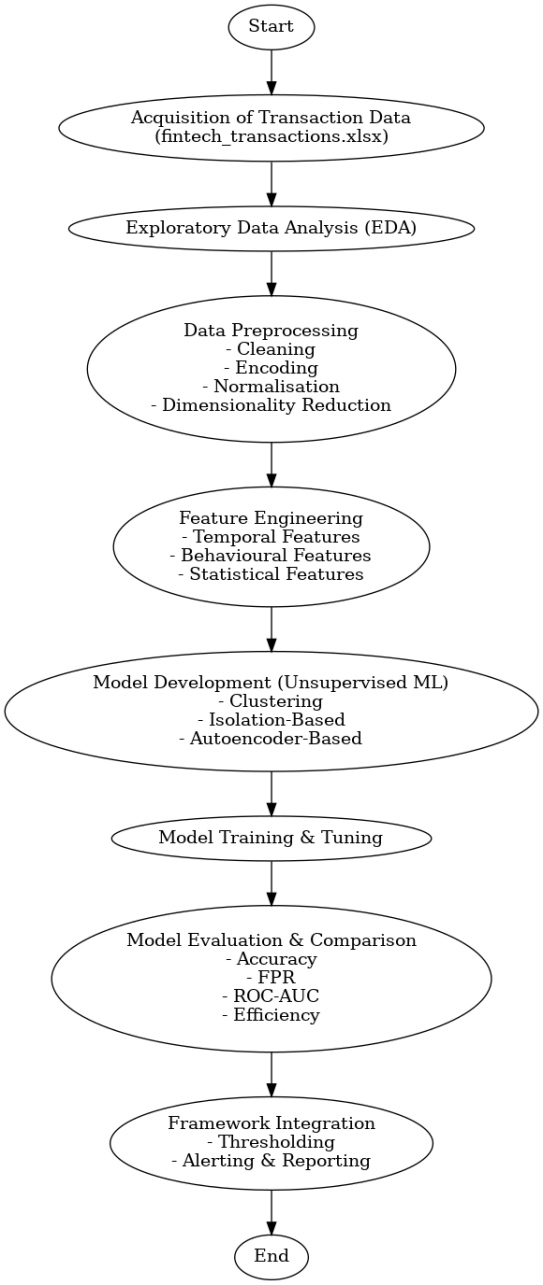

Figure 1: Flowchart of model framework

3.3 Data Preprocessing Pipeline

Data preprocessing was done on a Python based platform which preprocesses unstructured transaction data and transforms it into a structured format which can be used to construct a machine learning model. It consists of four steps.

i. Cleaning the Data: The first step of cleaning the data was changing the column containing the time stamp into a good datetime format such that the critical variables are numeric. Time stamp property has been exploited to produce other time properties like date, month and day of the week. Name of the categorical attributes were converted to lower-cases, and any additional space between the column names were eliminated. They have confirmed that the hour values were in the range of 0-23 with a flag variable being added to identify the outlier transactions. Binary attributes are normalized, duplications are eliminated and unique transaction IDs are checked. An indicator was proposed, in terms of outlier detection in terms of financial amounts, with the IQR methodology.

ii. Transforming the Features: A log-plus-one transformation was applied to the transaction amount to alleviate right-skewness.

The transformation is defined as: $\mathrm{amount}_{-\mathrm{log1p}} = \log(1 + \mathrm{amount})$

where log denotes the natural logarithm. The addition of 1 before taking the logarithm ensures numerical stability for zero-valued transactions. This is standard practice for financial transaction data, where amount distributions are heavily right-skewed due to the dominance of small-value transactions [7]. Behavioral indicators were also included e.g. is weekend to indicate transactions made on a Saturday or Sunday and is night to indicate transactions made in the late hours. To distinguish extremely large transactions, a high-amount flag was created by looking at the high IQR value to determine the high-amount transactions.

iii. Data Encoding: Categorical variables were turned into understandable dummy variables using one-hot encoding, whereas binary and flag features were not changed.

iv. Feature Selection and Dimensionality Reduction: A filter for variance thresholds removed features that had no variance, which reduced the feature set from 24 to 23 columns. that had no variance, which reduced the feature set from 24 to 23 columns. The mutual information ranking found the 23 most informative features and kept them. The most informative predictors of retention were found to be 23. Principal Component Analysis (PCA) was then used on the numeric subset, with two major components which together explained more than 95% of the total variance.

3.4 Feature Engineering

The feature engineering process added useful signals to the cleaned transactional dataset. These signals were then sorted into three groups: temporal, behavioral, and statistical. To stop data leakage, each user's features were calculated in a strict order of time.

There were temporal features like sine and cosine transforms of the hour of the day (hour_sin and hour_cos) to keep time cyclical, day-of-week and month indices to show seasonality, transaction recency (the number of hours since the user's last transaction), and rolling user transaction counts over seven days.

Behavioural features included device-user match rate (the frequency of prior device usage by the user), channel-switch and merchant-category-switch flags indicating changes relative to the preceding transaction, novel channel usage flags, deviation of the current transaction hour from the

user's prior average, and a rolling z-score of the current transaction amount relative to the user's own recent spending distribution.

Statistical features included amount z-scores within each country-merchant-category group, global and category-level amount percentile ranks, a winsorized residual measuring the excess amount beyond the 99th percentile cap, and an IQR-normalized amount within each merchant category.

3.5 Model Development

Three unsupervised model families were implemented based on their theoretical suitability and representativeness of distinct anomaly detection paradigms.

i. Clustering-Based Models (K-Means): K-Means was trained on the scaled and encoded feature set, with the anomaly score defined as the minimum Euclidean distance to any cluster centroid. The k values were selected based on domain reasoning and empirical elbow-method analysis during preliminary exploration. Specifically, k = 3 reflects a coarse three-way partitioning into broad behavioral profiles (low-value routine, moderate transfers, and high-value or irregular transactions) consistent with segmentation heuristics in fintech fraud literature [12]; k = 5 allows finer resolution of sub-segments such as daytime versus nighttime transaction clusters; and k = 8 was included to test whether higher granularity further improves anomaly discrimination. This range is consistent with comparable fintech anomaly detection studies [5].

   Values of k ∈ {3, 5, 8} were tested and the optimal configuration was selected based on validation Precision-Recall Area Under Curve (PR-AUC).

ii. Isolation-Based Models (Isolation Forest): Isolation Forest isolates anomalies through recursive

   random partitioning of the feature space. The hyperparameter ranges were defined as follows: n_estimators ∈ {100, 200, 300} because 100 trees is the accepted minimum for stable anomaly score estimates and values beyond 300 yield diminishing returns for datasets of this size [19], max_samples ∈ {256, 1024, 'auto'} to balance sub-sampling diversity, where 256 promotes faster isolation of extreme outliers, 1024 enables subtler deviations to be captured, and 'auto' serves as an adaptive baseline; and contamination fixed at 0.01 to reflect the observed fraud prevalence of approximately 0.9%, providing a principled threshold for converting continuous anomaly scores to binary labels during evaluation. Hyperparameters tuned included n_estimators ∈ {100, 200, 300} and max_samples ∈ {256, 1024, 'auto'}, with contamination fixed at 0.01 to reflect the observed fraud prevalence.

iii. Reconstruction-Based Models (PCA): PCA was used to learn a subspace capturing normal behaviour, with anomaly scores computed as mean squared reconstruction error. Variance thresholds of 90% and 95% were tested. The 90% threshold represents a widely used lower bound retaining sufficient variance to reconstruct normal transactions while discarding noise-driven components that might mask anomalies; the 95% threshold retains more dimensions and provides a stricter reconstruction baseline. The optimal threshold was selected based on validation PR-AUC. This range is consistent with PCA-based anomaly detection practice in financial datasets [21].

To get the highest PR-AUC, hyperparameters were tuned using a validation split of 20% of the training data. The performance of the test set was used to choose the final model. Data were split into two sets: a training set (80%) and a test set (20%) with stratification to preserve class balance.

3.6 Evaluation Metrics

Model evaluation was performed on a held-out test set comprising 10,000 transactions with a fraud prevalence of approximately 0.9%. Given the extreme class imbalance, traditional accuracy was not informative, and specialised metrics were employed.

ROC-AUC and PR-AUC (Average Precision, AP) were two threshold-independent metrics. ROC-AUC measured how well the model could tell the difference between fraud and non-fraud across all possible thresholds. PR-AUC measured the trade-off between precision and recall with a focus on the positive class.

On thresholds optimized by F1, the following calculations were carried out threshold-dependent measures such as Precision (TP/(TP+FP)), Recall (TP/(TP+FN)) and F1-score (the harmonic mean of precision and recall), specificity, false-positive rate, as well as Lift@θ (precision/fraud prevalence). Lift@θ is a measure of alert enrichment as compared to random selection.

On operational performance, Precision@K and Recall@K for the top K alerts (K $\in$ {50, 100, 200, 500}). This analysis replicates realistic review capacity limits, in that the number of alerts that can be examined in a single day is limited.

4. RESULTS AND DISCUSSION

4.1 Model Performance Overview

This section shows the real-world results of putting the suggested fraud detection framework into action. The main goal was to create and test unsupervised anomaly detection models that could find fake transactions in a very unbalanced financial dataset. Performance metrics for each model family: clustering-based (K-Means), isolation-based (Isolation Forest), and reconstruction-based (PCA) on both the training and test sets.

Table 1: Evaluation Metrics – Training Set

Model	ROC-AUC	PR-AUC	Threshold	TN	FP	FN	TP	Precision	Recall	F1	Lift@θ
Isolation Forest	0.6048	0.0139	−0.0251	38,050	1,590	333	27	0.0167	0.0750	0.0273	1.855
PCA (reconstruction)	0.6122	0.0130	0.0318	32,955	6,685	252	108	0.0159	0.3000	0.0302	1.767
K-Means (distance)	0.5424	0.0113	4.7183	39,204	436	352	8	0.0180	0.0222	0.0199	2.002

Table 2: Evaluation Metrics – Test Set

Model	ROC-AUC	PR-AUC	Threshold	TN	FP	FN	TP	Precision	Recall	F1	Lift@θ
Isolation Forest	0.5708	0.0146	−0.0386	9,495	415	84	6	0.0143	0.0667	0.0235	1.586
PCA (reconstruction)	0.5504	0.0099	0.0526	9,438	472	86	4	0.0084	0.0444	0.0142	0.934
K-Means (distance)	0.5269	0.0101	3.9425	9,011	899	80	10	0.0110	0.1111	0.0200	1.223

4.2 ROC Curve Analysis

FIGURE 2, demonstrates the aggregate receiver-operator characteristic curves, which provide a detailed comparison of the three paradigms of anomaly detection between the training and the test corpus.

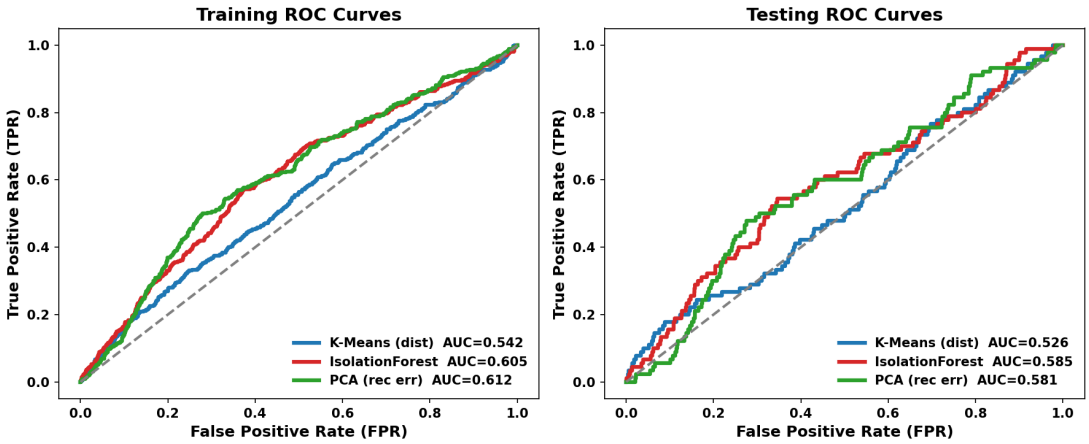

Figure 2: Combined ROC Curve Analysis for Training and Test Sets

On the training set, the ROC-AUC for Isolation Forest was about 0.60, the same for PCA reconstruction was about 0.61, and the AUC for K-Means was the lowest at about 0.54. Considering the test data that is not visible, Isolation Forester keeps its lead, scoring an AUC of about 0.57 and thus demonstrating strong generalisation. The next step would be PCA, which gives approximately 0.55, and K-Means is marginal at 0.53. The small training-to-testing AUC difference between Isolation Forest and PCA was also indicative of a weak over-fitting, as opposed to the continued poor performance of K-Means on both datasets.

4.3 Precision-Recall Analysis

FIGURE 3, below shows the combined Precision Recall curves that help in explaining how each analytical model balances precision and recall at a range of thresholds. Since fraud is a relatively uncommon incidence in the data, the Precision Recall view is especially edible; any significant improvements in initial accuracy can significantly reduce the burden on analysts. The Isolation

Forest and PCA models tend to produce more curves on the training data than K-Means, indicating a stronger concentration of genuine fraud cases within the transactions having the highest scores.

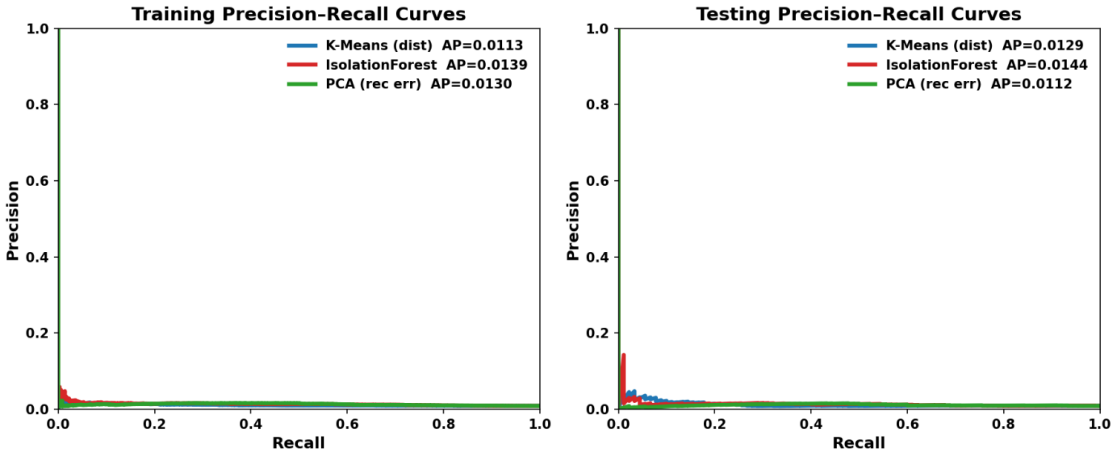

Figure 3: Precision-Recall Curves for Training and Test Sets

The Isolation Forest has the highest Average Precision (AP) in the test data, as well as, the most noticeable curve in the early-recall region, the very place where review teams operate. PCA has a slightly lower AP on testing than on training, which suggests that it is somewhat sensitive to subspace assumptions. On the other hand, the K-Means proves to be the weakest among the three in the test set, its initial-precision curve drops faster with higher recall.

4.4 Threshold Analysis

FIGURE 4, shows the operational performance of each model at their F1-tuned decision cut-off on the training and test data in the form of a visualization of the threshold (θ). In the training panel, Isolation Forest has the best balance between precision and the recall, thus the highest F1@θ. PCA is more likely to increase recall, but at the cost of reduced precision; K-Means, in its turn, has a more conservative attitude, therefore, displaying lower recall. Isolation Forest, again, in the test panel, has the most optimal total equilibrium with its accuracy holding a relatively high position and a recall that is substantially significant. The recall advantage that PCA has when training PCA does not apply when testing, as the precision decreases. K-Means is still the weakest of the three when it comes to testing. Its early precision curve drops faster as recall goes up.

4.5 Confusion Matrix Analysis

The observation of the true negatives as represented in the figure in the form of FIGURE 5, reveals that on the held-out test set, Isolation Forest has the highest number with about 9495 true negatives, closely followed by Principal Component Analysis with about 9438 and K-Means with about 9011. These findings imply that the Isolation Forest algorithm has a greater generalisation ability that maintains its strong expertise in appropriately identifying valid transactional behavior.

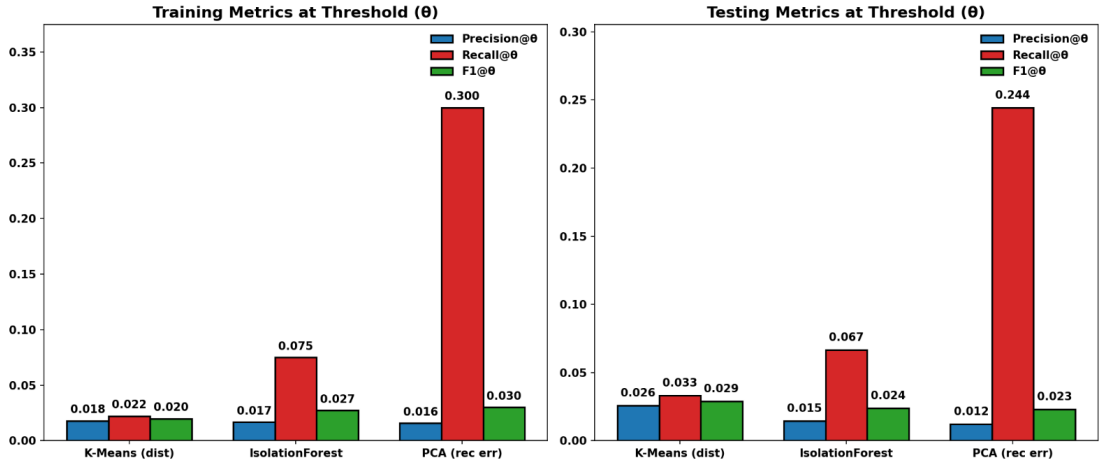

Figure 4: Threshold (θ) Analysis for Training and Test Sets

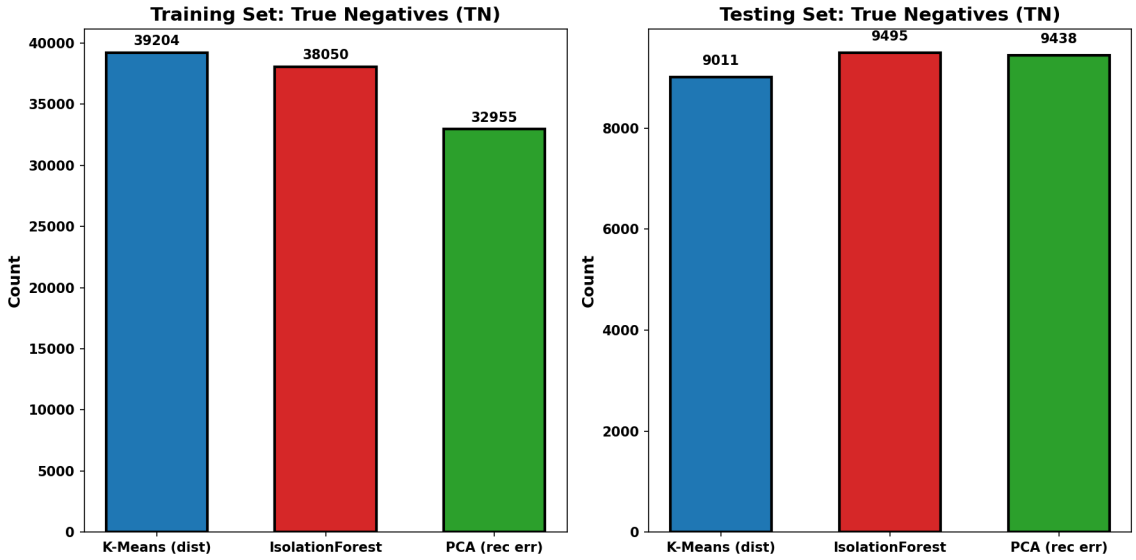

Figure 5: True Negatives (TN) by Model

The analysis of the false positives, which is given in FIGURE 6, shows that in the test partition, K-Means produces the highest number of false alarms (around 899), PCA produces 472, and Isolation Forest produces the lowest number (around 415). This trend shows that Isolation Forest is more precise and has a higher capacity to generalize. The increased false-positive rate of PCA is related to its bias towards recall, and Isolation Forest attains a more balanced trade-off.

FIGURE 7, presents the true positives (TP) for each model on the held-out test set. K-Means achieved the highest raw TP count (10), followed by Isolation Forest (6) and PCA (4). This result must be interpreted alongside the corresponding false-positive counts: K-Means attained its higher TP at the cost of 899 false positives and a precision of only 1.10%, whereas Isolation Forest produced fewer true positives but with substantially lower false alarms (415 FP) and higher precision (1.43%).

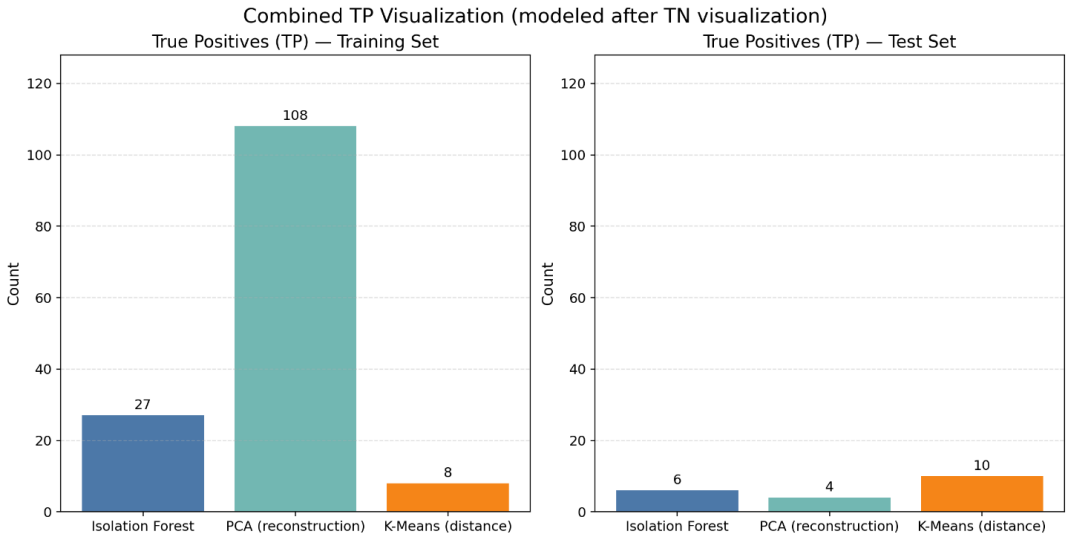

Figure 6: False Positives (FP) by Model

The relatively low absolute TP counts across all models are a direct consequence of the severe class imbalance (approximately 0.9% fraud prevalence) and the conservative threshold settings optimized for F1, and should not be interpreted in isolation from recall and precision metrics.

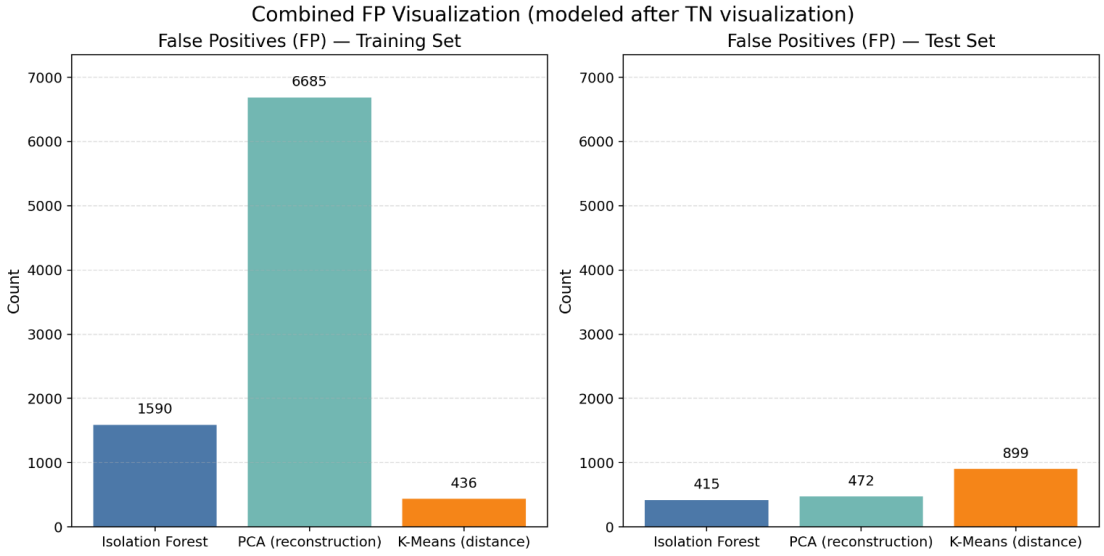

Figure 7: True Positives (TP) by Model

The analysis of the number of false-negatives, which are also presented in FIGURE 8, shows that the values of the test set are concentrated in a narrow range: Isolation Forest 84, PCA 86 and K-Means 80. The closeness of these values is an indicator of similar miss rates of the models when they are faced with invisible cases which highlights the overall difficulties of generalization in highly imbalanced data contexts.

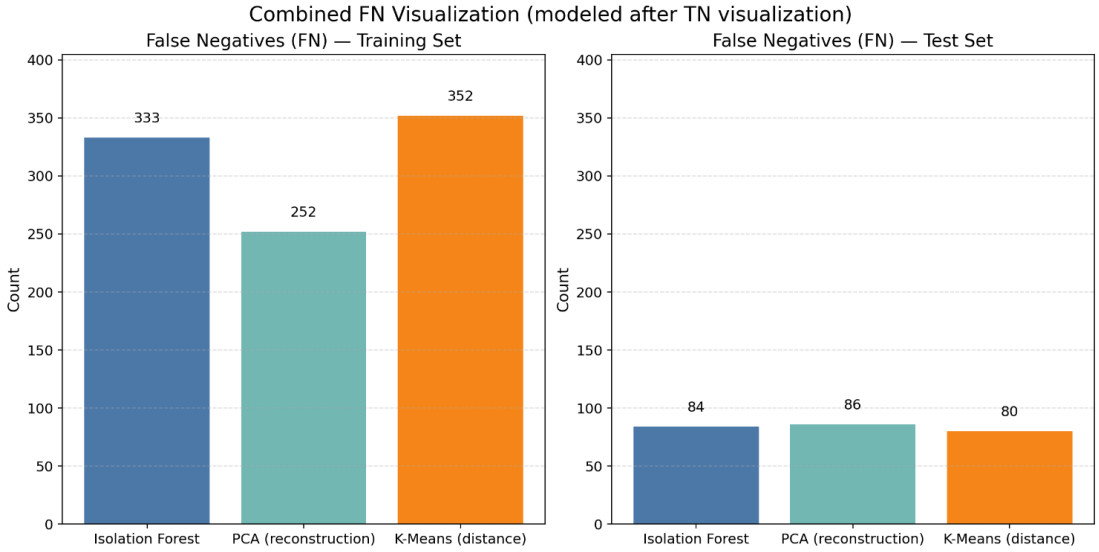

Figure 8: False Negatives (FN) by Model

4.6 Precision, Recall, and F1-Score

Precision metrics of all the investigated models were small, which is a result of the strong imbalance of classes in the data. However, distinguishable differences were noted. In particular, the Isolation Forest algorithm showed the best precision rate of 1.43% on the held-out test partition, outperforming those of K-Means (1.10%) and principal component analysis (0.84%), which in turn supports the fact that it has the best generalization ability on unseen observations at a relatively balanced detection rate (see FIGURE 9).

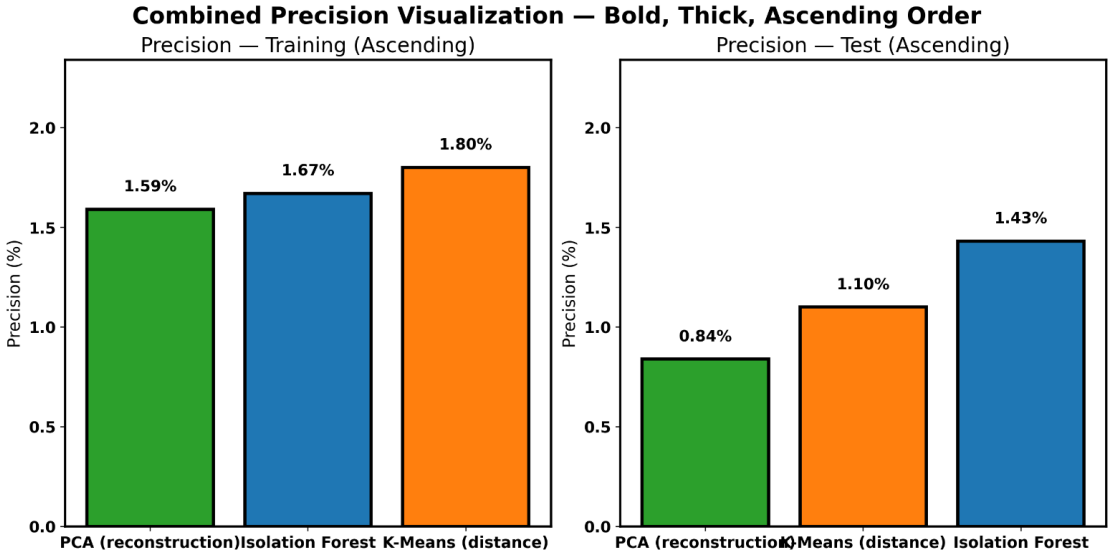

Figure 9: Precision by Model

FIGURE 10, shows that the K-Means clustering algorithm showed the largest recall rate of 11.11 percent, compared to the Isolation Forest (6.67 percent) and PCA (4.55 percent) on the same test partition. K-Means had a better raw recall score, but it was accompanied by a high false-positive load, indicating an inefficient procedure of generating alerts. Among fintech companies that are active in Nigeria, the acquisition of a balanced recall is especially relevant because the inability to identify the presence of fraud has a harmful effect on the stability of the business environment and the adherence to the rules imposed by the government.

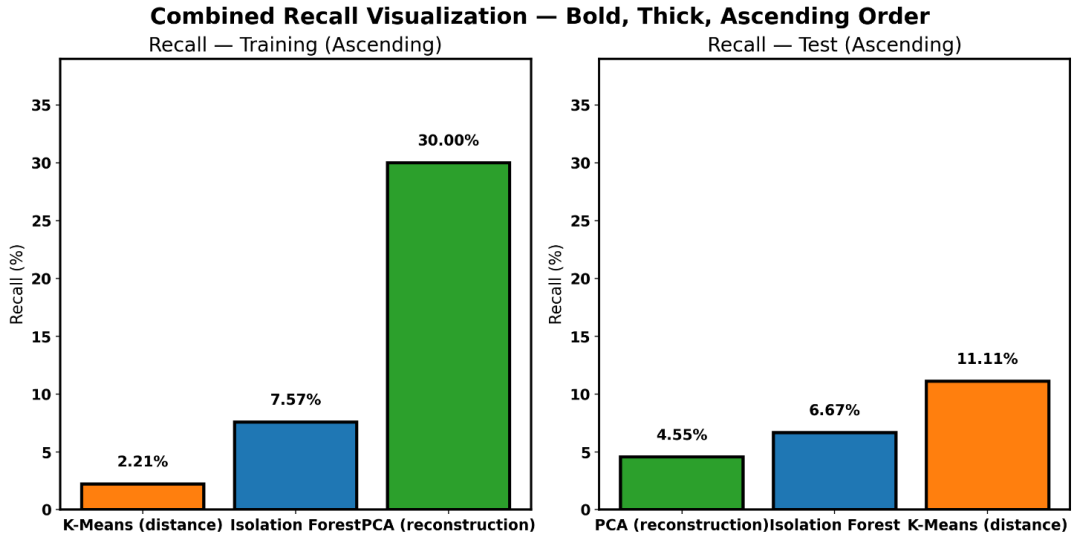

Figure 10: Recall by Model

The F-score analysis, as shown in FIGURE 11, indicates that, in the held-out test set, the Isolation Forest algorithm performed once again better with an F1-score of 2.38% compared to K-Means and PCA with the scores 1.96% and 1.52% respectively. This means that Isolation Forest has the best balance between precision and recall on data that hasn't been seen before. The relative simplicity of the F-scores in the models studied is aimed at revealing the inherent complexity of fraud detection in the situation with severe imbalance between the classes.

4.7 Lift@θ Analysis

Lift@θ is the enrichment of the true fraud cases in the flagged alerts compared with random selection. K-Means had the highest lift of approximately 2.002, Isolation Forest of approximately 1.855 and PCA of approximately 1.767 in the training set. Isolation Forest was also able to generalize better with a lift of about 1.586 on the test set, compared to K-Means (about 1.223) and PCA (about 0.934). It is important to note that the lift of PCA decreased to less than 1 with unseen data, which shows that its flagged alerts were less informative than a random selection, which highlights the danger of using too high an anomaly scoring threshold without calibration. The best trade-off between coverage and efficiency is Isolation Forest which retained a favourable lift on training and test sets.

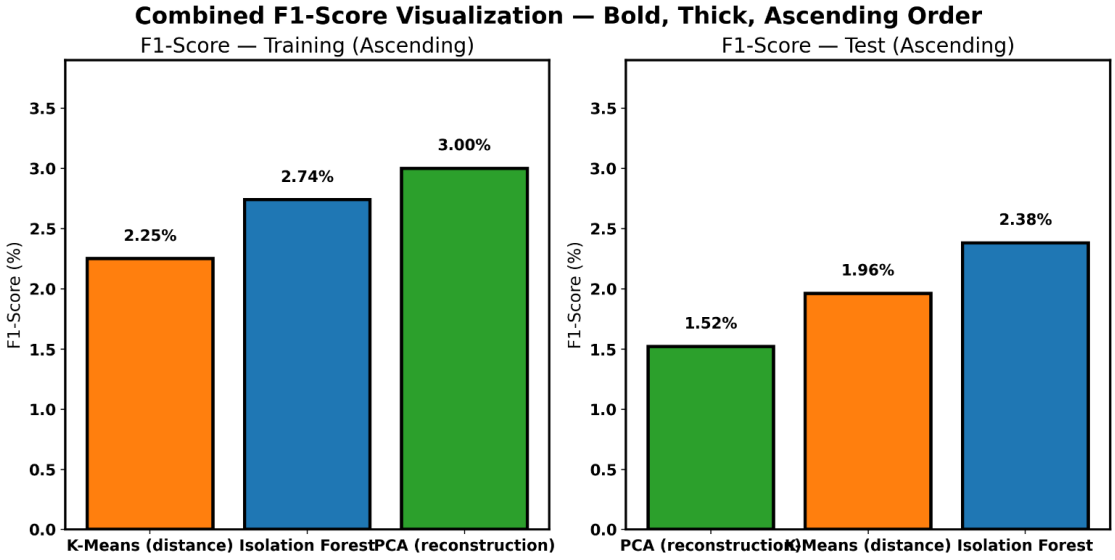

Figure 11: F1-Score by Model

5. CONCLUSION

This study proposed and tested the unsupervised machine learning models to identify anomalies in the context of the Nigerian fintech setting, dealing with the crucial issue of identifying fraudulent transfers in highly skewed datasets without the use of labelled data. Three unsupervised methods, which are Isolation Forest, PCA-based reconstruction, and K-Means clustering, were adopted and compared on real-world-inspired transactional data.

Isolation Forest was the most balanced in terms of metrics and reported moderate recall and precision and the highest Lift@θ on the test set (around 1.586), which shows that it enriched true fraud cases in flagged alerts strongly. PCA-based reconstruction had the best recall on the training set (30%) but had poor generalization with its Lift decreasing below 1 on the test set (about 0.934), and thus its alerts are not as informative as those of random selection. The most conservative and the best Lift was K-Means clustering, however, it had very low recall, which restricted coverage of fraud.

In general, these results highlight the trade-offs inherent in unsupervised anomaly detection: these models are optimized to achieve high recall, which increases the false positives, whereas conservative models minimize the false positives, which may deprive the fraud cases. The most viable method in the assessment of techniques is the Isolation Forest, which can be applicable to the case of Nigerian fintech startups since it has a favourable trade-off between the accuracy of detection and operational viability.

6. RECOMMENDATIONS

On the basis of the collected empirical data, the recommendations to both scholars and practicing professionals are as follows

(i) use the Isolation Forest algorithm as the default model of anomaly detection in the case of Nigerian fintech startups, because of its balance between precision and recall and its strong generalisation in different transaction settings

(ii) Introduce the dynamic threshold calibration processes in order to balance the trade-off between recall and precision to adapt the detection sensitivity to the operational risk appetites

(iii) To reduce the occurrence of false positives, unsupervised models should be combined with hybrid approaches, such as post-filtering mechanisms or simplified supervised re-scoring methods, at the same time increasing the interpretability of the model.

(iv) Incorporate lift -based monitoring metrics into real-time operational dashboards so as to enable ongoing assessment of the quality of alerts and to provide analytics stakeholders with performance metrics that may be acted upon

(v) Create regular retraining regimes with drift detection systems, which will allow the models to follow the constantly evolving patterns of transactions typical of the fintech ecosystem

(vi) Investigate explainability frameworks, e.g. SHAP or LIME, to provide feature-based information about the flagged anomalies and thus improve analyst credibility and guarantee regulatory expectations;

(vii) Take into account ensemble methods that combine Isolation Forest with reconstruction-based methods in situations where high recall is the most important

(viii) Invest in privacy-preserving, scalable infrastructure- adopt, e.g. federated learning, to enable a collaboration in detecting the presence of fraud across multiple fintech participants.

References

[1] Alonge EO, Eyo-Udo NL, Ubanadu BC, Daraojimba AI, Balogun ED, et al. Enhancing Data Security With Machine Learning: A Study on Fraud Detection Algorithms. J Front Multidiscip Res. 2021;2:19-31.

[2] Okare BP, Omolayo O, Aduloju TD. Designing Unified Compliance Intelligence Models for Scalable Risk Detection and Prevention in SME Financial Platforms. Int J Multidiscip Res Growth Eval. IJMRGE. 2024a;5:1421-1433.

[3] Shehadeh M. Exploring Cutting-Edge Algorithms in Fintech: Leveraging Machine Learning and Artificial Intelligence for Achieving Sustainability. In: From digital disruption to dominance. Bingley: Emerald Publishing Limited. 2025:225–244.

[4] Adewuyi A, Regina Nwangele C, Joyce Oladuji T, Oyeronke Akintobi A. Advances in Machine Learning for Credit Risk and Underwriting Automation: Emerging Trends in Financial Services. Int J Adv Multidiscip Res Stud. IJAMRS. 2023;3:1860-1877.

[5] Ahmad H, Kasasbeh B, Aldabaybah B, Rawashdeh E. Class Balancing Framework for Credit Card Fraud Detection Based on Clustering and Similarity-Based Selection (SBS). Int J Inf Technol. Springer Nature. 2023;15:325-333.

[6] Paul AA, Ogburie C. The Role of AI in Preventing Financial Fraud and Enhancing Compliance. GSC Adv Res Rev. 2025;22:269-282.

[7] Adejoh J, Owoh N, Ashawa M, Hosseinzadeh S, Shahrabi A, et al. An Adaptive Unsupervised Learning Approach for Credit Card Fraud Detection. Big Data Cogn Comput. 2025;9:217.

[8] John SA, Shonubi JA, Azuikpe PF, Ologun VO. Adoption of Ai-Driven Fraud Detection System in the Nigerian Banking Sector: An Analysis of Cost, Compliance, and Competency. 2025. arXiv preprint: https://arxiv.org/pdf/2511.00061.

[9] Ilori O. The Effectiveness of Data Mining in Detecting Financial Fraud: A Review and Applications. Int J Adv Multidiscip Res Stud. IJAMRS. 2025;5:1325-1339.

[10] Nwachukwu C, Akwiwu-Uzoma C, Ovuehor S, Durodola-Tunde K. Improved Machine Learning Algorithms for Fraud Detection in Fintech Companies. InInternational Conference on Financial Technology. Springer Nature. 2025:9-19.

[11] Salami IA, Popoola AD, Gbadebo MO, Kolo FH, Adesokan-Imran TO. Ai-Powered Behavioural Biometrics for Fraud Detection in Digital Banking: A Next-Generation Approach to Financial Cybersecurity. Asian J Res Comput Sci. 2025;18:473-494.

[12] Balciğoğlu, Y. S. Revolutionising risk management: AI and ML innovations in financial stability and fraud detection .IGI Global. 2024:109–138.

[13] Joyce Oladuji T, Oyeronke Akintobi A, Regina Nwangele C, Ajuwon A. A Model for Leveraging AI and Big Data to Predict and Mitigate Financial Risk in African Markets. Int J Adv Multidiscip Res Stud. IJAMRS. 2023;3:1843-1859.

[14] Ohiozua OM. Leveraging Financial Analytics for Fraud Mitigation and Maximising Investment Returns: A Comparative Analysis of the USA, Africa and Nigeria. Int J Res Publ Rev. 2024;5:1136-1152.

[15] Narayan M, Shukla P, Dileep Kumar M, Mani N. Generative AI in FinTech: Revolutionizing fraud detection, personalized advising, and predictive analytics. InGenerative AI in FinTech: Revolutionizing Finance Through Intelligent Algorithms. Springer Nature. 2025:137-154.

[16] Angela O, Atoyebi I, Soyele A, Ogunwobi E. Enhancing Fraud Detection and Prevention in Fintech: Big Data and Machine Learning Approaches. World J Adv Res Rev. 2024;24:2301-2319.

[17] Stojanović B, Božić J, Hofer-Schmitz K, Nahrgang K, Weber A, et al. Follow the Trail: Machine Learning for Fraud Detection in Fintech Applications. Sensors. 2021;21:1594.

[18] Adenuga AA, Gaffar O, Sikiru AO, Otunba M. Developing Comprehensive Ai-Driven Financial Technology Solutions That Democratise Access to Capital for Small Businesses. Int Sci Refereed Res J. 2023;6:320-343.

[19] Fonkem BN. Ai-Powered Risk Scoring Models for Real-Time Fraud Detection in Digital Banking Ecosystems. J Comput Anal Appl. JOCAAA. 2025;34.

[20] Ismail MM, Haq MA. Enhancing Enterprise Financial Fraud Detection Using Machine Learning. Eng Technol Appl Sci Res. 2024;14:14854-14861.

[21] Jiang S, Dong R, Wang J, Xia M. Credit Card Fraud Detection Based on Unsupervised Attentional Anomaly Detection Network. Systems. 2023;11:305.