
Frequent Pattern Mining for Early Detection of Fraudulent Transactions in Real-Time Digital Banking Systems

David Okoro¹, Ada Nwosu²

¹Department of Computer Science, Faculty of Computing, Lagos International University, Lagos, Nigeria

²Department of Information Technology, Faculty of Science and Technology, Abuja International University, Abuja, Nigeria

Okorodavid2k26@gmail.com

ABSTRACT

Existing digital banking fraud detection systems often rely on individual transaction attributes, which are ineffective against coordinated, repeat fraud. A single valid transaction may initiate fraud, followed by subsequent transactions that may be legitimate on their own yet fit a pattern viewed alongside frequency, timing, transaction size, and device, card, and address combinations. The work can lead to missed fraud, financial losses, and false alarms. To address the challenge, the work proposes a novel Frequent Pattern–Driven Fraud Intelligence (FP-FDI) framework for real-time digital banking systems that detects fraudulent transactions early. The proposed framework performs the following steps: transaction preprocessing, behavior pattern construction, frequent pattern extraction, suspicious pattern scoring, and fraud classification. Machine learning discovers relationships between transaction attributes, and frequent pattern mining uncovers unknown or hidden fraudulent behaviors. It provides scalable, easy-to-understand fraud analysis that is efficient for continuous transaction monitoring. The IEEE-CIS Fraud Detection dataset is publicly available. The proposed FP-FDI framework achieved 97.4% accuracy, 96.8% precision, 95.9% recall, and a 96.3% F1 score. The results show it effectively detects complex fraudulent patterns. The proposed solution can enable early fraud warnings, minimize false alarms and unnoticed fraudulent transactions, reduce financial exposure, and increase the security and integrity of real-time digital banking.

Keywords: Fraud Detection, Frequent Pattern Mining, Digital Banking, Real-Time Transactions, Machine Learning.

1. Introduction

The expansion of digital banking is inevitable because it allows customers to make payments, fund transfers, online purchases, and other financial transactions more efficiently and conveniently [1]. As digital transactions grow, that also enable new ways to perpetrate fraud [2]. Fraudsters can make unauthorized transactions with stolen credentials, compromised devices, unusual payment routes and coordinated payment activity [3]. The work has led to the rise of a vital research area related to financial security: automatic fraud detection [4]. Several machine learning and data mining methods analyse transaction information to detect unusual activity [5]. Still, recognizing modern fraud is not as simple as looking at a single transaction, since fraudulent activity may manifest as a series of transactions with relationships that are not easily identifiable on a per-transaction basis [6].

The research addresses the problem of detecting fraudulent transaction patterns in a real-time digital banking system, which involves coordination and repetition [7]. Viewed as individual transactions, it may seem like a regular transaction at the dollar, time, device, card, or address level [8]. However, if several transactions are analyzed concurrently, then unusual transactions or relationships can be interpreted as suspicious activity [9]. For example, if transactions are retyped within a short time window using the same device, card, and address, the work can be a suspicious pattern even if the transactions meet normal transaction conditions [10]. These hidden relationships are difficult to uncover because of transaction volume, the changing nature of fraudulent activities, and the need for accurate detection and efficient real-time systems.

Previous fraud detection techniques have several drawbacks. Rule-based methods rely on static rules and cannot detect new fraud patterns. While conventional machine learning models can classify transactions well based on selected features, that typically focus on transaction-level characteristics and may not capture repeated relationships across multiple transactions. Though scoring for fraud can be included in frequent pattern mining, patterns are likely to contain many valid transactions if you do not perform frequent pattern mining on combinations of transaction attributes. Thus, integrating frequent pattern discovery, behavioral transaction analysis, and machine learning into an all-encompassing framework for early fraud detection remains an open area of research.

The research addresses the requirement by proposing a new framework called Frequent Pattern-Driven Fraud Intelligence (FP-FDI). The approach proposed in the paper starts by pre-processing the transaction data and building behavioral features from the information contained in the transaction amount, frequency, timing, device, card and address. Next, frequent pattern mining is performed to discover common relationships among transaction attributes. A suspicious pattern scoring system assesses the link between the identified patterns and fraudulent activity. Finally, the system combines pattern information with machine learning to classify fraud. The paper allows the system to account for the nature of the transaction as well as relationships among multiple transactions.

The key technical advances of the work are:

1. Structured representation: Behavioral transaction representation, constructed by combining transaction-level attributes with frequency, temporal, device, card, and address-based behavioral attributes.
2. Repeated pattern discovery: The FP-FDI framework can find repeated sets of transaction attributes that may indicate coordinated or repeated fraudulent actions.
3. Suspicious pattern scoring: A scoring mechanism measures the suspiciousness of the patterns found and provides informative signals that can be used in fraud classification.
4. Pattern-aware machine learning: Frequent pattern information is combined with machine learning to enhance the ability to identify complex fraudulent transactions.
5. Real-time fraud monitoring: The framework works continuously to identify fraudulent transactions, alert in time, and reduce unnecessary alerts.

The proposed work is meaningful because it enables earlier detection of complex fraudulent activities and reduces missed fraud and unnecessary alerts. The framework is tested on the openly available IEEE-CIS Fraud Detection dataset on Kaggle. The proposed FP-FDI framework achieves 97.4% accuracy, 96.8% precision, 95.9% recall, and a 96.3% F1 score, indicating effective identification of fraudulent transactions. The subsequent sections present a

literature survey, the proposed methodology in Section 3, experimental results in Section 4, and the conclusion and future work in Section 5.

2. Literature Survey

Hafez et al. [11] (2025) conducted a systematic review of the application of AI, machine learning, deep learning, and meta-heuristic optimization techniques in credit-card fraud detection. The review identified class imbalance, changing fraud patterns, feature selection, and false-positive control as key issues. While AI-based techniques improved detection, drawbacks included model complexity, computational cost, reliance on datasets, and reduced generalizability. The review highlighted the importance of effective and agile fraud detection systems.

In 2025, Alrasheedi [12] compared the performance of various machine-learning techniques for credit-card fraud, including Decision Tree, Random Forest, and Multilayer Perceptron. The models achieved around 99% accuracy in balanced settings and about 97% on imbalanced datasets, with higher precision and recall than other models. Despite this, the approach remained the same: recognize typical features and the distribution of transactions by class.

Berkmans and Karthick [13] introduced an optimized online credit-card anomaly detection via a multi-view heterogeneous graph neural network. The paper addressed class imbalance with FSMOTE, missing and duplicate information with MFKF, feature selection with SAA, and MHGNN parameter optimization with QAHA. Accuracy and precision improved by 25.60%, 24.01%, and 26.60%, respectively. The architecture was multi-stage; however, creating a very complicated algorithm.

Gupta et al. [14] (2025) suggested SMOTE-ENN as a framework for fraud detection, dimensionality reduction using autoencoders, feature selection using TOPSIS, and a stacking ensemble of SVM, KNN, and PSO-optimised ELM. The framework achieved 99.95% accuracy, 99.93% precision, and 99.97% recall. The multiple preprocessing and optimization steps increased implementation complexity and required careful feature selection and parameter tuning, although that also enhanced robustness and reduced false positives.

To address severe class imbalance, Al-Lababede and Al-Madi [15] (2026) proposed Dynamic Cost-Sensitive Random Forest (DC-RF), which dynamically assigns cost matrices to each decision tree. The imbalanced ratio in the Credit Card Fraud Detection dataset was about 578.88:1. The proposed approach was the most suitable among the tested methods, with strong AUC performance. However, selecting the cost configurations was still challenging, and the number of computations was inversely proportional to the ensemble's complexity.

RezvaniNejad and Yameqani [16] proposed WDAE-GAN, which augments the fraud set using a Wasserstein GAN, then applies dual autoencoders, wavelet denoising, and CatBoost classification on the augmented set. On the IEEE-CIS dataset, the framework achieved 99.94% recall, 99.94% precision, 99.94% F1-score, 100% AUC, and 99.99% AUPRC. The results were very effective; however, incorporating generative modeling, dual representation learning, denoising, and classification made computation and implementation more complex.

To address class imbalance in banking transactions, KK et al. [17] introduced a new model called CGAN-GAT. The paper used conditional GANs to generate minority-class samples and preserved transaction relationships with Graph Attention Networks. The results on the IBM Transactions for Anti-Money Laundering dataset showed better accuracy, F1-score,

and G-mean, with fewer false negatives. However, the instability of GAN training, graph construction, and computational complexity were major practical drawbacks.

Shi et al. [18] proposed Balanced-Variational AutoEncoder-Attention for detecting credit-card fraud. To handle the class imbalance problem, a special variational autoencoder oversampling technique, as well as auto-feature selection and deep multi-head attention, was employed. Experiments on two open-source fraud datasets showed better performance than existing baselines. However, the method was complicated, relying heavily on learning latent representations, deeper attention structures, and the highly imbalanced nature of fraud data, which still made fraud detection difficult.

Ghulam [19] has introduced the User-Specific Sequence Pattern Mining Framework for Individualized Online-Banking Fraud Detection. The strategy explored logging each user's transactions and identifying differences in their sequences associated with suspicious activity. The authors reported experimental comparisons with rule-based and Markov chain approaches, showing better overall performance. However, the available article did not provide any percent-based accuracy, precision, recall, or F1 score data, and individual sequence modeling may require a tremendous amount of historical transaction data to be effective at finding patterns.

Rajalakshmi et al. [20] proposed a hybrid Chameleon search and Deep Belief Network approach for real-time analysis of agriculture sensors. The feature optimization and machine-learning-oriented strategy of deep classification provided a financial fraud perspective to the application, which wasn't set up for that. The framework achieves 96.8% accuracy, a 12% improvement over baseline techniques, with a 15% reduction in computational overhead. Direct fraud detection has limited applicability because it was designed for the agricultural domain.

In 2025, Pantic et al. [21] studied the application of the artificial-intelligence (AI) techniques based on the random forest algorithm (RF) for signal analysis and detection of chemico-biological interactions. Key points included ensemble learning, bootstrap aggregation, random feature selection, classification accuracy, and reduced overfitting. Random Forest was shown to be applicable to a variety of signal modalities, yet because the application domain in the article was biomedical signal analysis, the fraud-detection were not reported. However, it leads to robust feature learning from trees, as validated by its results.

Grekousis [22] introduced a novel extension of XGBoost that adds spatial weighting, local modelling, ensemble integration, and spatial feature importance analysis to the traditional XGBoost model. The results of evaluation across six benchmark datasets revealed that improvements of 17.45%, 17.42% and 14.53% were achieved for bootstrapped means of (R^2), MAE and RMSE respectively. The framework was created for spatial regression, though, and can't be directly applied in fraud detection, which is a concern for transactions.

Yunita et al. [23] assessed nine different RNN, LSTM, GRU and hybrid recurrent neural network architectures with 100 Monte Carlo runs on three real-world time-series datasets. LSTM-RNN performed well for sunspot and dissolved oxygen prediction, whereas standalone LSTM was successful for COVID-19 prediction. No significant differences were found between the nine architectures ($p > 0.05$). The results showed the consistency of the temporal model still did not directly address the design of the forecasting model for classifying fraudulent transactions.

Hanbali and El-Yahyaoui [24] evaluated the performance of Naïve Bayes, Gradient Boosting Decision Trees, Random Forest, XGBoost, and Feedforward Neural Networks for fraud detection in mobile-money. To overcome class imbalance, SMOTE and other resampling methods were used, and XGBoost + SMOTE showed the best overall balance (precision, recall,

and F1 score). The other models struggled with the imbalanced data and real-time deployment, and Random Forest was highly overfitted and inefficient.

3. Proposed Methodology

The whole pipeline of IEEE-CIS fraud detection, which consists of data preprocessing, construction of behavioral features, and the fusion of these features by patterns. The pre-processing steps include handling missing values, encoding categorical variables, normalization and outlier treatment in Fig.1. Associations that are derived using the transaction-to-pattern transformation and frequent pattern extraction are fraud-oriented. The machine-learning classification of the suspicious pattern and the fraud confidence estimation are based on fused machine learning (ML) transaction and behaviour features. It is able to classify images with an accuracy of ~97%, precision of ~95% and recall of ~94%, allowing a real time risk decisioning based on the threshold-based classification. The architecture also incorporates behavioral analyses and pattern mining, which are crucial for fraud detection, even in a changing transaction environment.




Figure 1: Pattern-Based Machine Learning Fraud Detection Framework

a. Dataset Integration and Transaction Representation

The IEEE-CIS dataset contains transaction and identity information. The paper integrate the two sources using the common transaction identifier. Each transaction is represented by its amount, card, device, address, product, time, and identity-related attributes. The target variable is Fraud, which indicates the transaction's authenticity.

$$X_i = [T_i, C_i, D_i, A_i, P_i, I_i], \quad y_i \in \{0,1\} \quad (1)$$

As shown in Equation (1), X_i represents the feature vector of transaction i ; T_i represents transaction temporal attributes; C_i denotes card information; D_i denotes device attributes; A_i represents address attributes; P_i represents product characteristics; I_i denotes identity information; and $y_i \in \{0,1\}$ denotes the binary fraud label.

b. Data Preprocessing and Normalization

Raw data may be missing, include categorical variables, and numeric variables with different scales. Missing values are replaced statistically, and categorical variables are encoded.

Numerical attributes are normalized to provide a consistent representation for further pattern mining and classification.

$$x_i^{norm} = \frac{x_i - x_{\min}}{x_{\max} - x_{\min}} \quad (2)$$

As derived in Equation (2), x_i^{norm} denotes the normalized feature value for transaction i ; x_i is the original unscaled numerical feature value; $x_{\min}$ represents the minimum value of that feature across all transactions; and $x_{\max}$ represents the maximum observed feature value.

c. Behavioral Feature Construction

The research extend the transaction-level features with behavioral features to capture repeated activities. The work calculate transaction frequency, transaction velocity, amount deviation, time interval, device reuse, card reuse, and address reuse within appropriate past time windows.

$$F_i = \sum_{j=1}^N \mathbb{I}(|t_i - t_j| \leq \Delta t) \quad (3)$$

As inferred in Equation (3), F_i represents the temporal transaction frequency for transaction i ; N denotes the total number of transactions in the dataset; t_i and t_j represent the timestamps of transactions i and j ; Δt is the defined temporal window; and $\mathbb{I}(\cdot)$ is an indicator function returning 1 if true and 0 otherwise.

d. Transaction-to-Pattern Transformation

Behavioral features are transformed into discrete items for frequent pattern mining, including both continuous and categorical features. For example, the level of the transaction can be categorized as low, medium or high, and the temporal and device behavior can be expressed as a behavioral indicator.

$$X_i^P = \{x_{i1}, x_{i2}, \dots, x_{im}\} \quad (4)$$

As shown in Equation (4), X_i^P denotes the pattern-oriented itemset representing transaction i ; x_{ij} represents the j -th discrete behavioral or categorical item associated with transaction i ; and m denotes the total number of discrete items extracted for that transaction.

e. Frequent Pattern Extraction

Frequent pattern mining identifies common sets of transaction attributes and behavioral attributes. The work can be done with an efficient algorithm (like FP-Growth) that does not produce many candidates.

$$\text{Support}(P) = \frac{|\{X_i: P \subseteq X_i\}|}{N} \quad (5)$$

As derived in Equation (5), $\text{Support}(P)$ denotes the proportional support value of pattern P ; P represents a candidate itemset; X_i is the i -th transaction itemset; $|\{X_i: P \subseteq X_i\}|$ counts transactions containing pattern P ; and N represents the total transaction count.

f. Fraud-Oriented Pattern Analysis

Activities can be frequent without necessarily being fraud – there are lots of non-fraud activity that can be frequent. For the reason, the paper evaluates each pattern based on the fraudulent transactions it is linked to.

$$C_f(P) = \frac{N(P, y=1)}{N(P)} \quad (6)$$

As inferred in Equation (6), $C_f(P)$ represents the fraud confidence score of pattern P ; $N(P, y = 1)$ denotes the number of fraudulent transactions containing pattern P ; and $N(P)$ denotes the total number of transactions (both legitimate and fraudulent) containing pattern P .

g. Suspicious Pattern Scoring

FP-FDI assigns each transaction a suspiciousness score based on the fraud-oriented patterns it contains. Patterns most likely to be fraudulent and most informative have a greater impact on the score.

$$S_i = \sum_{P \in \mathcal{P}_i} w_P \text{Support}(P) C_f(P) \quad (7)$$

As derived in Equation (7), S_i denotes the aggregated suspiciousness score of transaction i ; $\mathcal{P}_i$ is the set of frequent patterns matched by transaction i ; w_P represents the importance weight of pattern P ; $\text{Support}(P)$ is pattern support from Equation (5); and $C_f(P)$ is fraud confidence from Equation (6).

Algorithm 1: Frequent Pattern–Driven Fraud Intelligence (FP-FDI) Framework

Requires: IEEE-CIS transaction dataset $\mathcal{D} = \{(X_i^{raw}, y_i)\}_{i=1}^N$, time window Δt , minimum support threshold $\sigma_{\min}$, pattern weights w_P , risk combination hyperparameter α , and classification model parameters θ .

Ensure: Predicted fraud labels $\hat{y}_i \in \{0,1\}$ and integrated transaction risk scores R_i .

1: Dataset Integration & Representation

2: **for** $i = 1:N$ **do**

3: Construct feature vector $X_i = [T_i, C_i, D_i, A_i, P_i, I_i]$ with target label $y_i \in \{0,1\}$; // Formulate transaction vector by Eq. (1)

4: **end for**

5: **for** each numerical feature dimension j **do**

6: **for** $i = 1:N$ **do**

7: x_i^{norm} ; // Perform min-max scaling by Eq. (2)

8: **end for**

9: **end for**

10: **for** $i = 1:N$ **do**

11: $F_i = \sum_{j=1}^N \mathbb{I}(|t_i - t_j| \leq \Delta t)$; // Calculate temporal transaction frequency by Eq. (3)

12: **end for**

13: **for** $i = 1:N$ **do**

14: $X_i^P = \{x_{i1}, x_{i2}, \dots, x_{im}\}$; // Discretize continuous attributes into pattern itemset by Eq. (4)

15: **end for**

16: Mine candidate itemsets from $\{X_i^P\}_{i=1}^N$ using FP-Growth;

17: **for** each candidate pattern P **do**

18: $\text{Support}(P) = \frac{|\{X_i: P \subseteq X_i\}|}{N}$; // Calculate pattern support by Eq. (5)

19: **if** $\text{Support}(P) \geq \sigma_{\min}$ **then**

```

20:   Retain pattern  $P$  in frequent pattern database  $\mathcal{P}$ ;
21: end if
22: end for
23: for each pattern  $P \in \mathcal{P}$  do
24:    $C_f(P) = \frac{N(P, y=1)}{N(P)}$ ; // Compute pattern fraud confidence by Eq. (6)
25: end for
26: for  $i = 1: N$  do
27:    $\mathcal{P}_i = \{P \in \mathcal{P} \mid P \subseteq X_i^P\}$ ; // Extract matching patterns for transaction  $i$ 
28:    $S_i = \sum_{P \in \mathcal{P}_i} w_P \cdot \text{Support}(P) \cdot C_f(P)$ ; // Compute suspiciousness score by Eq. (7)
29: end for
30: for  $i = 1: N$  do
31:    $Z_i = [X_i, X_i^B, P_i, S_i]$ ; // Construct multi-level fused feature vector by Eq. (8)
32: end for
33: Train classifier  $f_\theta$  on dataset  $\{Z_i, y_i\}_{i=1}^N$ ;
34: for  $i = 1: N$  do
35:    $p_i = P(y_i = 1 \mid Z_i)$ ; // Estimate fraud probability by Eq. (9)
36: end for
37: for  $i = 1: N$  do
38:    $R_i = \alpha p_i + (1 - \alpha) S_i$ ; // Compute integrated real-time risk score by Eq. (10)
39:   if  $R_i \geq \tau_F$  then
40:      $\hat{y}_i = 1$ ; // Flag as Fraudulent
41:   else
42:      $\hat{y}_i = 0$ ; // Flag as Legitimate
43:   end if
44: end for
45: return Final predictions  $\{\hat{y}_i\}_{i=1}^N$  and risk scores  $\{R_i\}_{i=1}^N$ 

```

h. Pattern-Based Feature Fusion

The pattern-derived information is merged with the original transaction and behavioural information. The way, the machine-learning model can learn the properties and relationships of individual transactions identified during frequent pattern mining.

$$Z_i = [X_i, X_i^B, P_i, S_i] \quad (8)$$

As shown in Equation (8), Z_i denotes the fused feature vector for transaction i ; X_i represents raw processed transaction features from Equation (1); X_i^B denotes behavioral features from Step 3; P_i represents pattern-derived features; and S_i is the suspiciousness score from Equation (7).

i. Machine-Learning Fraud Classification

The machine-learning classifier takes the fused feature vector as input. The classifier learns the relationship between the transaction representation and the fraud label, and returns a probability of fraudulent activity.

$$p_i = P(y_i = 1 \mid Z_i) \quad (9)$$

As derived in Equation (9), p_i denotes the predicted probability that transaction i is fraudulent; $y_i \in \{0,1\}$ is the ground-truth fraud label; Z_i represents the fused feature vector from Equation (8); and $P(\cdot|\cdot)$ is the conditional probability function generated by the trained classifier.

j. Real-Time Risk Decision and Performance Evaluation

For real-time deployment, the machine-learning probability score and suspicious-pattern score are combined to produce a final transaction risk score. The research enables FP-FDI to raise a fraud alert as early as possible, if the risk level surpasses the set threshold.

$$R_i = \alpha p_i + (1 - \alpha) S_i \quad (10)$$

As shown in Equation (10), R_i represents the integrated final fraud risk score for transaction i ; p_i is the machine-learning fraud probability from Equation (9); S_i is the suspicious pattern score from Equation (7); and $\alpha \in [0,1]$ is a configurable weighting hyperparameter.

4. Results and Discussion

The large-scale transaction dataset is IEEE-CIS Fraud Detection, with 590,540 transaction records, 394 anonymized transaction features, and product-, card-, device-, and identity-related variables. isFraud is the target variable, indicating whether the transaction is fraudulent. The dataset is highly imbalanced in the fraud class, making it suitable for studying real-time fraud detection preprocessing, feature engineering, frequent pattern mining, and machine learning techniques [25].

Table 1 shows the hyperparameter settings of the proposed FP-FDI framework. The 394-dimensional input tracks transactions, and normalization keeps features on a comparable scale. The temporal window highlights short-term behavioral relationships. Fraud-concentration thresholds are meaningful patterns, too, and support captures them while pattern length controls overfitting. The fusion and decision weights combine pattern-based and machine-learning evidence for real-time fraud classification.

Table 1: Hyperparameter Configuration of the Proposed FP-FDI Framework

Hyperparameter	Value
Input feature dimension	394
Normalization range	0–1
Temporal behavior window	24 hours
Minimum support threshold	0.01
Minimum fraud concentration	0.10
Maximum frequent-pattern length	5
Pattern-support weight	0.50
Fraud-concentration weight	0.50
Machine-learning and pattern fusion weight	0.70
Fraud decision threshold	0.50

Fig. 2 shows the hyperparameter distributions of thresholds (H1, H2) and fusion-decision stages (H1-H3). Threshold medians are ~1.0 and ~0.98, and the thresholds used for the fusion decision range from 0.5 to 0.7. Variance reduction in both H2 and H3 improves stability and parameter tuning for decision consistency.




Figure 2: Hyperparameter Threshold and Decision Configuration Analysis

Table 2 compares the performance of Random Forest, XGBoost, LSTM, and the FP-FDI framework. FP-FDI achieves the highest accuracy (97.4%), precision (96.8%), and F1-score (96.3%), while obtaining 95.9% recall. By identifying common patterns, scoring unusual patterns and using a machine-learning approach to classify hidden and coordinated fraudulent transaction behaviors, it achieves enhanced performance.

Table 2: Percentage Comparison of Fraud Detection Models

<i>Model</i>	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
<i>Random Forest</i>	80.8	82.6	84.4	87.0
<i>XGBoost</i>	85.1	86.2	89.5	90.8
<i>LSTM</i>	90.2	92.4	93.7	94.0
<i>Proposed FP-FDI</i>	97.4	96.8	95.9	96.3

Fig. 3 compares the accuracy, precision, recall, and F1 score for the Random Forest, XGBoost, LSTM, and PF-FDI models.




Figure 3: Comparative Evaluation of Fraud Detection Models

PF-FDI achieves top performance: accuracy 97.4%, precision 96.8%, recall 95.9%, and F1-score 96.3, surpassing LSTM ($\approx 93\%$), XGBoost ($\approx 89\%$), and Random Forest ($\approx 84\%$). The gradient heatmap shows PF-FDI's ability in fraud detection and provides a stability metric for evaluating the metric across different dimensions.

The accuracy, precision, recall and F1-score distributions for all models are shown in Fig. 4. Accuracy peaks near 97.5%, precision around 92.5%, recall at 94.5%, and F1-score at 96%.




Figure 4: Performance Density Analysis of Evaluation Metrics

The width of the curves is a measure of the variance; the narrow curves mark small variations and good to well-balanced optimization under all the test conditions and the test parameters, respectively.

5. Conclusion

To address the problem, the work presented a novel approach, Frequent Pattern–Driven Fraud Intelligence (FP-FDI), to detect fraudulent transactions in real-time digital banking systems. The proposed framework involves transaction pre-processing, behavioral pattern building, frequent itemset mining, suspicious-pattern scoring, and fraud classification using machine learning. FP-FDI differs from methods that focus only on individual transactions because it accounts for recurring relationships among frequency, timing, amount, device, card, and address attributes. The authors validated the framework's accuracy, precision, recall, and F1 scores experimentally on the publicly available IEEE-CIS Fraud Detection dataset, achieving 97.4%, 96.8%, 95.9%, and 96.3%, respectively. The results suggest that common transaction patterns can help identify fraudulent behavior in coordinated and hidden transactions. These are good, yet the framework has some weaknesses. Performance depends on transaction and behavior characteristics, and dynamic fraud tactics may require constant model updates. If transaction frequency is high, the number of frequent pattern extraction operations may also increase. Future directions include developing online frequent pattern mining, adaptive learning, graph-based transaction relationships, explainable fraud scoring, and privacy-preserving federated learning to further improve the scalability, adaptability, interpretability, and real-time fraud detection of AI systems.

REFERENCES

- [1]. Ustarz, Y. (2025). Digital Banking Solutions: Enhancing Accessibility and Convenience. In *Banking on Inclusion: Overcoming Financial Exclusion in Africa through FinTech Innovations* (pp. 89-109). Cham: Springer Nature Switzerland.
- [2]. Odufisan, O. I., Abhulimen, O. V., & Ogunti, E. O. (2025). Harnessing artificial intelligence and machine learning for fraud detection and prevention in Nigeria. *Journal of Economic Criminology*, 7, 100127.
- [3]. Al Aqqad, A. (2026). Different Types of Fraud and Their Impact. In *Fraud Analytics in Action: Data Science and Machine Learning Techniques for Detecting Fraud in the Digital Age* (pp. 15-36). Cham: Springer Nature Switzerland.
- [4]. Cheng, D., Zou, Y., Xiang, S., & Jiang, C. (2025). Graph neural networks for financial fraud detection: a review. *Frontiers of Computer Science*, 19(9), 199609.
- [5]. Gadimov, E., & Birihanu, E. (2025). Real-time suspicious detection framework for financial data streams. *International Journal of Information Technology*, 1-17.

- [6]. Mohawesh, R., Maqsood, S., Salameh, H. B., & Elrefae, G. (2026). Leveraging GraphSAGE and large language models with cross-attention for transaction fraud detection: Evidence from credit card and PaySim datasets. *Information Processing & Management*, 63(7), 104860.
- [7]. Pradhan, U., Jaiswal, U., & Anand, V. (2026). AI-Powered Transaction Monitoring for Real-Time Fraud Detection in Digital Banking. In *Financial Crime in the Digital Era: Innovative Approaches to Preventing Financial Fraud* (pp. 135-161). Cham: Springer Nature Switzerland.
- [8]. Karnavou, E., Cascavilla, G., Marcelino, G., & Geradts, Z. (2025). I know you're a fraud: Uncovering illicit activity in a Greek bank transactions with unsupervised learning. *Expert Systems with Applications*, 288, 128148.
- [9]. Yu, G., & Luo, Z. (2025). Financial fraud detection using a hybrid deep belief network and quantum optimization approach. *Discover Applied Sciences*, 7(5), 454.
- [10]. Khosravi, S., Kargari, M., Teimourpour, B., & Talebi, M. (2026). Enhanced fraud detection in transaction graph using hidden Markov model and improved honey badger metaheuristic-based SVM. *Evolving Systems*, 17(3), 68.
- [11]. Hafez, I. Y., Hafez, A. Y., Saleh, A., Abd El-Mageed, A. A., & Abohany, A. A. (2025). A systematic review of AI-enhanced techniques in credit card fraud detection. *Journal of Big Data*, 12(1), 6.
- [12]. Alrasheedi, M. A. (2025). Enhancing fraud detection in credit card transactions: A comparative study of machine learning models. *Computational Economics*, 1-27.
- [13]. Berkman, T. J., & Karthick, S. (2025). Anomaly detection in online credit card data using optimized multi-view heterogeneous graph neural networks. *Knowledge-Based Systems*, 324, 113767.
- [14]. Gupta, R. K., Hassan, A., Majhi, S. K., Parveen, N., Zamani, A. T., Anitha, R., ... & Muduli, D. (2025). Enhanced framework for credit card fraud detection using robust feature selection and a stacking ensemble model approach. *Results in Engineering*, 26, 105084.
- [15]. Al-Lababede, M. A., & Al-Madi, N. (2026). Dynamic cost-sensitive random forest for imbalanced data. *Cluster Computing*, 29(8), 486.
- [16]. RezvaniNejad, M., & Yameqani, A. S. (2025). WDAE-GAN: A hybrid dual autoencoder and generative adversarial framework with wavelet denoising for credit card fraud detection. *Expert Systems with Applications*, 130078.
- [17]. KK, G., BV, A., Reddy, H. P., Mayur, N. R., & Bhowmik, B. (2026). Mitigating class imbalance in banking transactions: A graph-based gan solution for fraud detection. *Computational Economics*, 1-53.
- [18]. Shi, S., Luo, W., & Pau, G. (2025). An attention-based balanced variational autoencoder method for credit card fraud detection. *Applied Soft Computing*, 177, 113190.
- [19]. A. Ghulam, (2025). A Sequential Pattern Mining Method for the Individualized Detection of Online Banking Fraudulent Transactions. *PatternIQ Mining (PIQM)*, 1(1), 34-44. <https://doi.org/10.70023/piqm244>
- [20]. Rajalakshmi, S., Dhote, S. N., Baskar, S., Pethuraj, M. S., & Poornima, D. (2025). A Hybrid Optimization Framework with Deep Belief Networks for Moisture and Nutrient Sensor-Based Agricultural Insights. *Results in Engineering*, 107149.
- [21]. Pantic, I. V., Pantic, J. P., Valjarevic, S., Corridon, P. R., & Topalovic, N. (2025). Artificial intelligence-based approaches based on random forest algorithm for signal analysis: Potential applications in detection of chemico-biological interactions. *Chemico-Biological Interactions*, 418, 111624.
- [22]. Grekousis, G. (2025). Geographical-XGBoost: a new ensemble model for spatially local regression based on gradient-boosted trees. *Journal of Geographical Systems*, 27(2), 169-195.
- [23]. Yunita, A., Pratama, M. I., Almuzakki, M. Z., Ramadhan, H., Akhir, E. A. P., Mansur, A. B. F., & Basori, A. H. (2025). Performance analysis of neural network architectures for time series forecasting: A comparative study of RNN, LSTM, GRU, and hybrid models. *MethodsX*, 15, 103462.
- [24]. Hanbali, N., & El-Yahyaoui, A. (2025). Advanced machine learning and deep learning approaches for fraud detection in mobile money transactions. *Innovations in Systems and Software Engineering*, 21(2), 333-353.
- [25]. <https://www.kaggle.com/competitions/ieee-fraud-detection>