2026 Volume: 02 Issue: 11 www.ijrpa.com ISSN 2456-9995 Research Article

International Journal Research Publication Analysis

Page: 01-17

AI-BASED FINANCIAL FRAUD DETECTION PLATFORM

G. Sathwika*[1], Y. Tulasi[1], B. Jayachandra[1], G. Shivaram[1], Dr. P. Srihari[1]

[1]Department of CSE (Artificial Intelligence & Data Science) MRIT – Deemed to be University Rajam, Andhra Pradesh, India.

Article Received: 09 October 2026
Article Revised: 29 October 2026
Published on: 18 November 2026

***Corresponding Author: G. Sathwika**
Department of CSE (Artificial Intelligence & Data Science) MRIT – Deemed to be University Rajam, Andhra Pradesh, India.

DOI: https://doi-doi.org/101555/ijrpa.3203

ABSTRACT: The growth of online banking and digital payment services has increased exposure to financial fraud, while static rule-based detection systems adapt poorly to new fraud patterns. This paper presents an AI-based financial fraud detection platform that combines a supervised classifier (XGBoost) with an unsupervised anomaly detector (Isolation Forest). Each transaction is described by behavioural features such as amount, transaction velocity, amount deviation, time since the previous transaction, location anomaly and merchant risk. The outputs of both models are combined to assign a Low, Medium or High risk category, SHAP (SHapley Additive exPlanations) is used to explain individual predictions, and high-risk transactions generate fraud alerts that are delivered to dashboards through WebSockets. The platform is implemented with a React.js frontend, a FastAPI backend and an SQLite database accessed through SQLAlchemy, with JWT-based authentication, role-based access control for Customer, Analyst and Admin users, and PDF audit reports generated with ReportLab; 160 backend Pytest cases passed. On a synthetic dataset of 5,000 transactions (4,750 legitimate and 250 fraudulent), the hybrid model achieved a precision of 97.92%, recall of 94.00%, F1-score of 95.92%, ROC-AUC of 99.98% and PR-AUC of 99.62%. These results were obtained on synthetic data only and do not establish equivalent performance on real banking or UPI transactions.

KEYWORDS: Financial fraud detection, XGBoost, Isolation Forest, hybrid machine learning, explainable AI, SHAP, risk classification.

1. INTRODUCTION

Online banking, the Unified Payments Interface (UPI) and mobile payment applications have made money transfers faster and more convenient, but they have also created opportunities

for fraud, including unauthorised fund transfers, identity theft, phishing and account misuse. Such incidents cause financial losses and reduce trust in digital payment systems, so detecting fraudulent transactions accurately and early is an important requirement for financial institutions and payment platforms.

Traditional fraud detection systems rely on predefined rules and fixed conditions. They can identify known fraud patterns, but they are less effective against new or changing techniques and may raise false alerts or miss unusual transactions that do not match any rule. Machine learning (ML) offers an alternative: models can learn from historical transaction data and recognise relationships between transaction characteristics, and anomaly detection methods can flag deviations from normal behaviour.

This paper describes an AI-based financial fraud detection platform built on a hybrid ML approach. XGBoost, a gradient-boosted decision tree algorithm, learns fraud patterns from labelled transactions, while Isolation Forest identifies transactions that are unusual relative to normal behaviour without needing labels. The platform derives behavioural features from each transaction, assigns a Low, Medium or High risk category, explains predictions with SHAP, and raises alerts for high-risk transactions. It is delivered as a web application with Customer, Analyst and Admin roles, transaction monitoring, dashboards and PDF audit reports. The contribution is the integration of established techniques into one working platform; no new learning algorithm is proposed.

The remainder of the paper is organised as follows. Section 2 reviews related work, Section 3 identifies the research gap, Section 4 states the contributions, Section 5 describes the methodology, Section 6 presents the system implementation, Section 7 reports the evaluation on synthetic data, and Sections 8 and 9 give the conclusion and future scope.

2. LITERATURE SURVEY / RELATED WORK

Ten recent studies were reviewed (Table 1). They fall into four groups: hybrid supervised–unsupervised frameworks, supervised payment-fraud models, AI-based risk assessment and monitoring studies, and survey articles.

Hybrid frameworks. Monteiro et al. [1] combined Isolation Forest and XGBoost for electricity fraud detection on real utility data, using SMOTETomek for class imbalance and SHAP for interpretation; they reported an AUC-ROC of 0.999 and an F1-score of 0.77, with 5-fold cross-validation. Shanaa and Abdallah [8] combined supervised classification (XGBoost) with unsupervised anomaly detection for credit card fraud and reported an accuracy of 99.98%, precision of 95.69%, recall of 92.50% and F1-score of 94.07%.

Chapwanya and Gorejena [9] combined PCA, XGBoost, Local Outlier Factor and Isolation Forest for insurance fraud and reported 92% accuracy, 92% precision and 96% recall, with fewer false positives than the approaches they compared against.

Supervised payment-fraud models. Dalal et al. [10] applied an XGBoost model with hyperparameter optimisation to financial payment fraud and evaluated it with precision, recall, F1-score, accuracy and 10-fold cross-validation. Their approach is supervised and depends on labelled data.

AI-based risk assessment and monitoring. Battu et al. [2] studied AI, ML, deep learning and graph neural networks for real-time fraud detection and risk assessment, reporting detection accuracy above 95%, and noted high data and computational requirements and limited interpretability. Wang et al. [3] applied AI and data mining to financial fraud detection, emphasising data preprocessing, feature engineering, real-time monitoring and early warning. Elalkaoui et al. [5] reviewed the integration of big data and AI in financial decision-making, covering hybrid and ensemble models, scalability, concept drift and explainability.

Surveys. Uddin [4] reviewed AI methods for preventing financial crime, including supervised, unsupervised, hybrid and explainable approaches. Lucas and Jurgovsky [6] surveyed ML for credit card fraud detection, covering class imbalance, concept drift and sequential transaction behaviour. Btoush et al. [7] systematically reviewed 181 papers on credit card fraud detection with ML and deep learning, and identified the need for better handling of imbalanced data, fewer false positives and further work on hybrid approaches.

Table 1: Summary of the reviewed studies.

Ref.	Domain / focus	Year	Technique	Reported findings
[1]	Electricity fraud	2025	Isolation Forest + XGBoost, SHAP, SMOTETomek	AUC-ROC 0.999; F1-score 0.77; 5-fold CV
[2]	FinTech risk and fraud	2026	AI/ML, deep learning, GNN, behavioural analysis	Accuracy > 95%
[3]	Financial fraud (data mining)	2024	Data mining, feature engineering, real-time monitoring	Evaluates detection, monitoring and early warning
[4]	Financial crime (review)	2026	Supervised, unsupervised, hybrid, XAI	Review; common metrics discussed
[5]	Big data and AI in finance (review)	2026	ML, deep learning, hybrid and ensemble models	Review; accuracy, RMSE, MAPE, RZ considered
[6]	Credit card fraud (survey)	2020	Supervised and unsupervised ML; imbalance; concept drift	Survey; precision, recall, F1, AUC discussed
[7]	Credit card fraud (systematic review)	2023	ML and DL, including Isolation Forest	Review of 181 papers

[8]	Credit card fraud	2025	Hybrid supervised + unsupervised, XGBoost	Accuracy 99.98%; precision 95.69%; recall 92.50%; F1 94.07%
Ref.	**Domain / focus**	**Year**	**Technique**	**Reported findings**
[9]	Insurance fraud	2025	PCA, XGBoost, LOF, Isolation Forest	Accuracy 92%; precision 92%; recall 96%
[10]	Financial payment fraud	2022	Hyperparameter-tuned XGBoost	Precision, recall, F1, accuracy; 10-fold CV

3. RESEARCH GAP

The reviewed studies show that XGBoost is widely used for classification on labelled transactions and that Isolation Forest is useful for finding unusual patterns. They also leave several gaps:

1. Many systems depend mainly on supervised learning and may not identify previously unseen fraud patterns [7], [10].
2. Several approaches target a single domain, such as electricity, credit card or insurance fraud, rather than digital financial transactions in general [1], [8], [9].
3. Supervised classification and unsupervised anomaly detection need to be combined to cover both known and unusual fraud [7].
4. Transaction behaviour, including frequency, amount variation, time gaps and location changes, needs to be part of the analysis.
5. Predictions should be explainable so that users can see why a transaction was flagged [4].
6. Few studies describe an integrated platform that combines fraud prediction, risk classification, alert management, real-time notification and audit reporting; the reviewed papers [4], [6], [7] are surveys or model-level studies rather than complete monitoring platforms.

4. NOVELTY AND CONTRIBUTIONS

To address these gaps, the work integrates established techniques into one platform rather than introducing a new learning algorithm. The main contributions are:

• A hybrid fraud detection approach that combines XGBoost classification with Isolation Forest anomaly detection on the same behavioural transaction features, addressing gaps 1 to 4.

• Low, Medium and High risk classification of every transaction, with SHAP-based explanations of the factors behind a prediction (gap 5).

• A fraud alert module that stores alerts with a status and delivers real-time notifications through WebSockets, together with role-based access for Customer, Analyst and Admin users, transaction monitoring and PDF audit report generation (gap 6).

• An evaluation of the hybrid model on a synthetic dataset of 5,000 transactions, with the limitations of synthetic evaluation stated explicitly.

5. PROPOSED METHODOLOGY

5.1 Overall System Architecture

The platform is organised in five layers (Fig. 1). The **user interface layer** (React) lets customers, analysts and administrators submit transactions, monitor activity, view dashboards and review alerts. The **backend processing layer** (FastAPI) handles authentication, transaction processing, feature extraction, fraud prediction, alert management and report generation. The **machine learning layer** contains the XGBoost and Isolation Forest models whose outputs feed the hybrid risk assessment. The **explainability and alert layer** uses SHAP to explain predictions and generates alerts according to the assigned risk category, with WebSockets used for real-time delivery. The **database and reporting layer** uses SQLAlchemy with SQLite to store application records, and ReportLab to produce PDF audit reports.

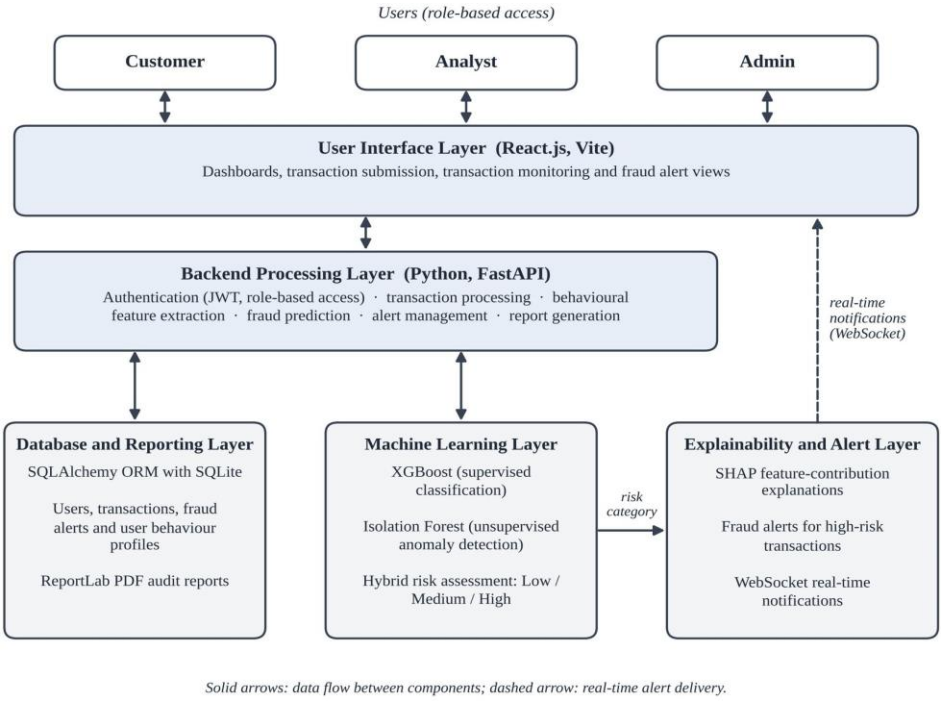

Fig. 1: Overall architecture of the proposed platform, showing the five layers and the flow of data between them.

5.2 Dataset Description

The models were trained and evaluated on a synthetically generated transaction dataset of 5,000 records, of which 4,750 are legitimate and 250 are fraudulent (5% fraud). Each record carries transaction attributes such as amount, transaction frequency, time between transactions, location information and merchant risk. XGBoost uses the labelled records for supervised classification, and Isolation Forest uses the same features to identify unusual patterns. Synthetic data allows controlled model development without confidential financial records, but it may not reproduce the complexity of real transaction behaviour (Section 7.4).

5.3 Data Preprocessing

Before modelling, the transaction records are organised and cleaned, missing values are handled, and the data are converted into a format suitable for the models. The workflow diagrams (Figs. 2 and 3) additionally list normalisation and encoding among the generic preprocessing stages. The prepared records are then passed to the feature engineering stage.

5.4 Feature Engineering

Each transaction is represented by six behavioural features (Table 2). Taken together they describe how large a transaction is, how frequently the account has been used, how far the amount departs from the user's usual behaviour, how recently the previous transaction occurred, and whether the location or merchant is unusual. The transaction submission form also records account-related fields, namely account age and the number of previous transactions; these are not among the model input features listed in the project documentation and are not treated as such here.

Table 2: Transaction features used by the fraud detection models.

No.	Feature	Contribution to fraud detection
1	Transaction amount	Unusually large amounts are a common indicator of fraudulent transfers.
2	Transaction velocity	Reflects how frequently transactions occur; a burst of transactions in a short period may indicate account misuse.
3	Amount deviation	Measures how far the amount differs from the user's normal transaction behaviour.
4	Time since previous transaction	Very short gaps between transactions can indicate automated or rapid misuse.
5	Location anomaly	Flags transactions made from locations that differ from the user's usual pattern.
6	Merchant risk	Reflects the risk level associated with the merchant involved.

5.5 XGBoost-Based Fraud Detection

XGBoost (Extreme Gradient Boosting) is a supervised algorithm that builds decision trees

sequentially, each new tree attempting to reduce the errors of the previous ones, and combines their outputs into a final prediction. In the platform it is trained on labelled transactions, and for each new transaction it outputs a fraud probability from which the transaction is classified as legitimate or fraudulent. It is suitable for this task because it handles structured tabular data and can learn non-linear relationships between features. The workflow is shown in Fig. 2.

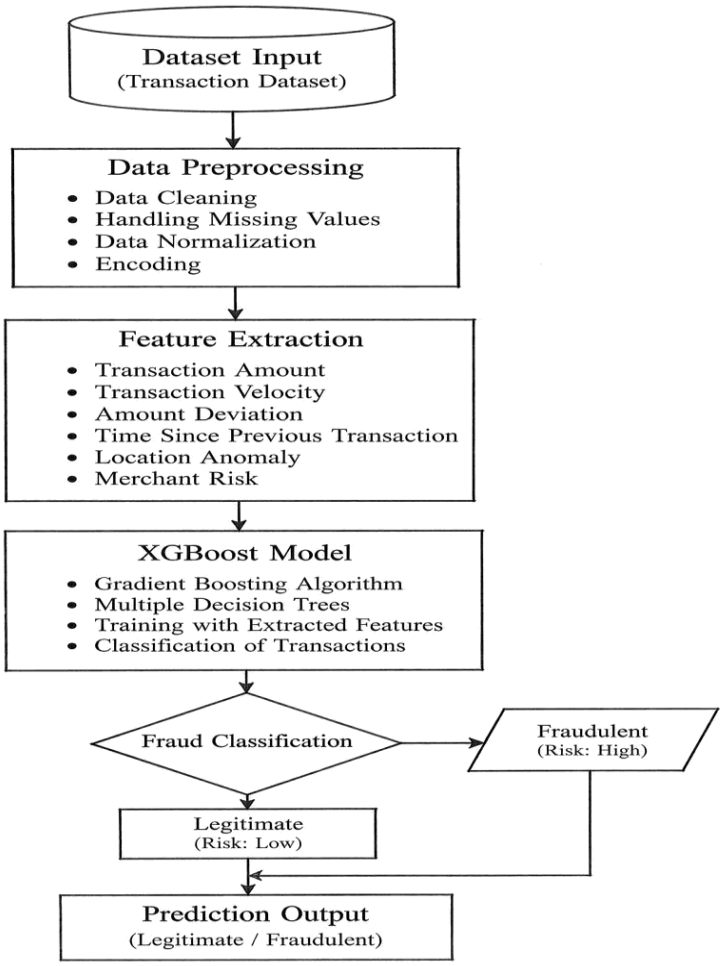

Fig. 2: Workflow of the XGBoost-based fraud classification module.

5.6 Isolation Forest-Based Anomaly Detection

Isolation Forest is an unsupervised algorithm that isolates observations by repeatedly selecting a random feature and a random split value. Anomalous points are separated in fewer splits, so a shorter average path length across the isolation trees gives a higher anomaly score. In the platform the same behavioural features are supplied to Isolation Forest, and its anomaly score identifies transactions that differ from normal behaviour even if they do not resemble any labelled fraud pattern. The workflow is shown in Fig. 3.

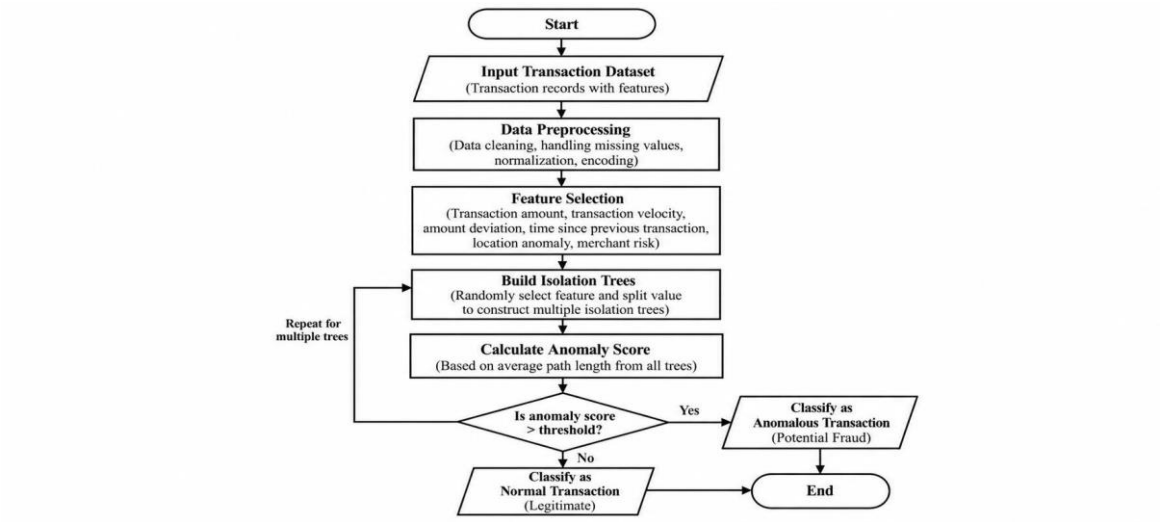

Fig. 3: Workflow of the Isolation Forest-based anomaly detection module.

5.7 Hybrid Fraud Detection Approach

The two models play complementary roles. XGBoost recognises fraud patterns that were present in the labelled training data, whereas Isolation Forest highlights unusual behaviour that may not have been seen before. As shown in Fig. 4, both models receive the same engineered features and run independently, XGBoost producing a fraud probability and Isolation Forest an anomaly score. The two outputs are then combined in a hybrid risk assessment step, which produces the risk category used for alerts and monitoring. This allows the platform to flag suspicious transactions that either model alone might miss. The combination is described here at the workflow level.

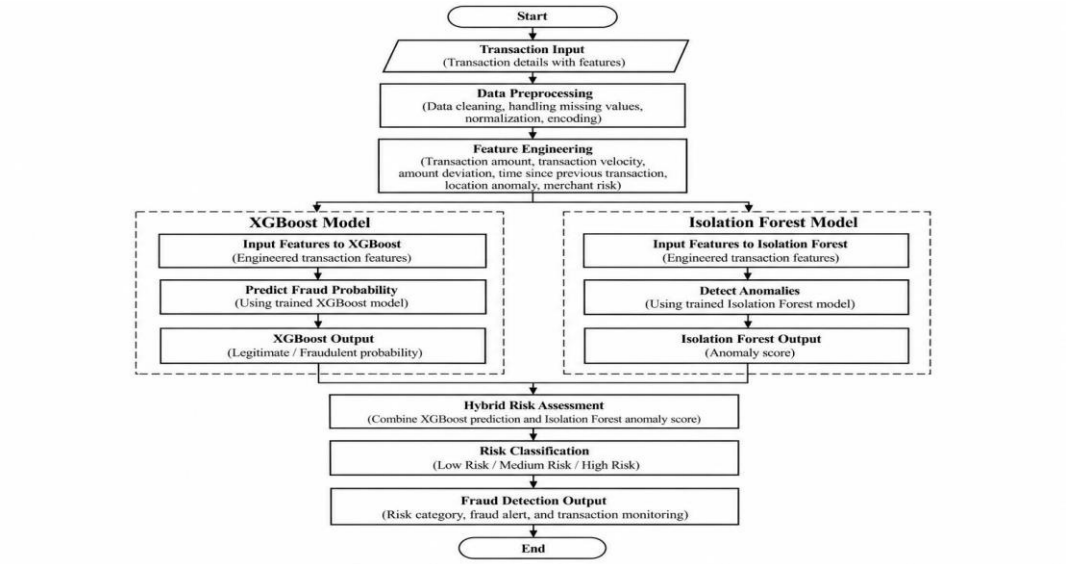

Fig. 4: Workflow of the hybrid fraud detection model combining XGBoost and Isolation Forest outputs.

5.8 Risk Assessment and Classification

The hybrid assessment assigns each transaction to one of three categories (Fig. 5). Low-risk transactions are treated as normal and generate no alert. Medium-risk transactions are flagged for review, with additional monitoring or verification. High-risk transactions are treated as potentially fraudulent and generate a fraud alert for the concerned analysts and users. The category boundaries are defined by a lower and an upper threshold on the combined risk score. The platform flags and reports suspicious transactions for review; it is not described as blocking payments automatically.

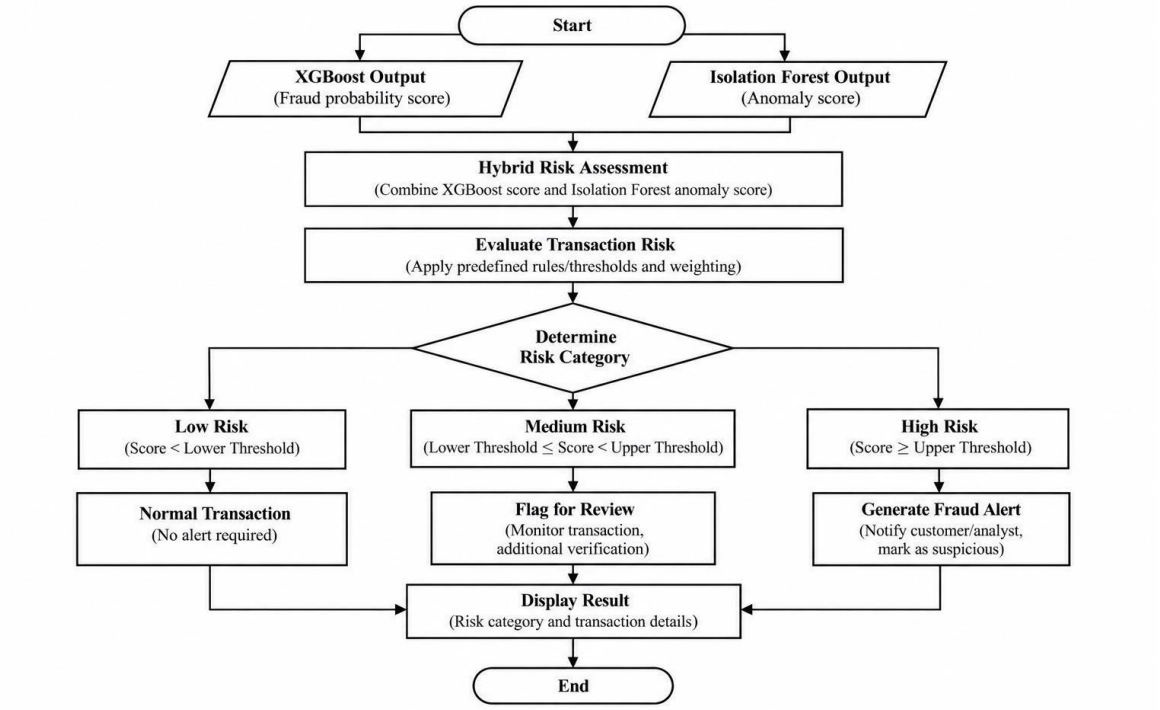

Fig. 5: Workflow of risk assessment and Low/Medium/High classification.

5.9 SHAP-Based Explainable AI

To make predictions understandable, the platform uses SHAP, which assigns each input feature a Shapley value representing its contribution to a particular prediction. Features with positive values increase the predicted fraud risk and features with negative values reduce it. For a flagged transaction, the contributions of the features are converted into an explanation that identifies the main factors behind the prediction (Fig. 6). This improves the transparency of the process and supports the investigation of suspicious activity.

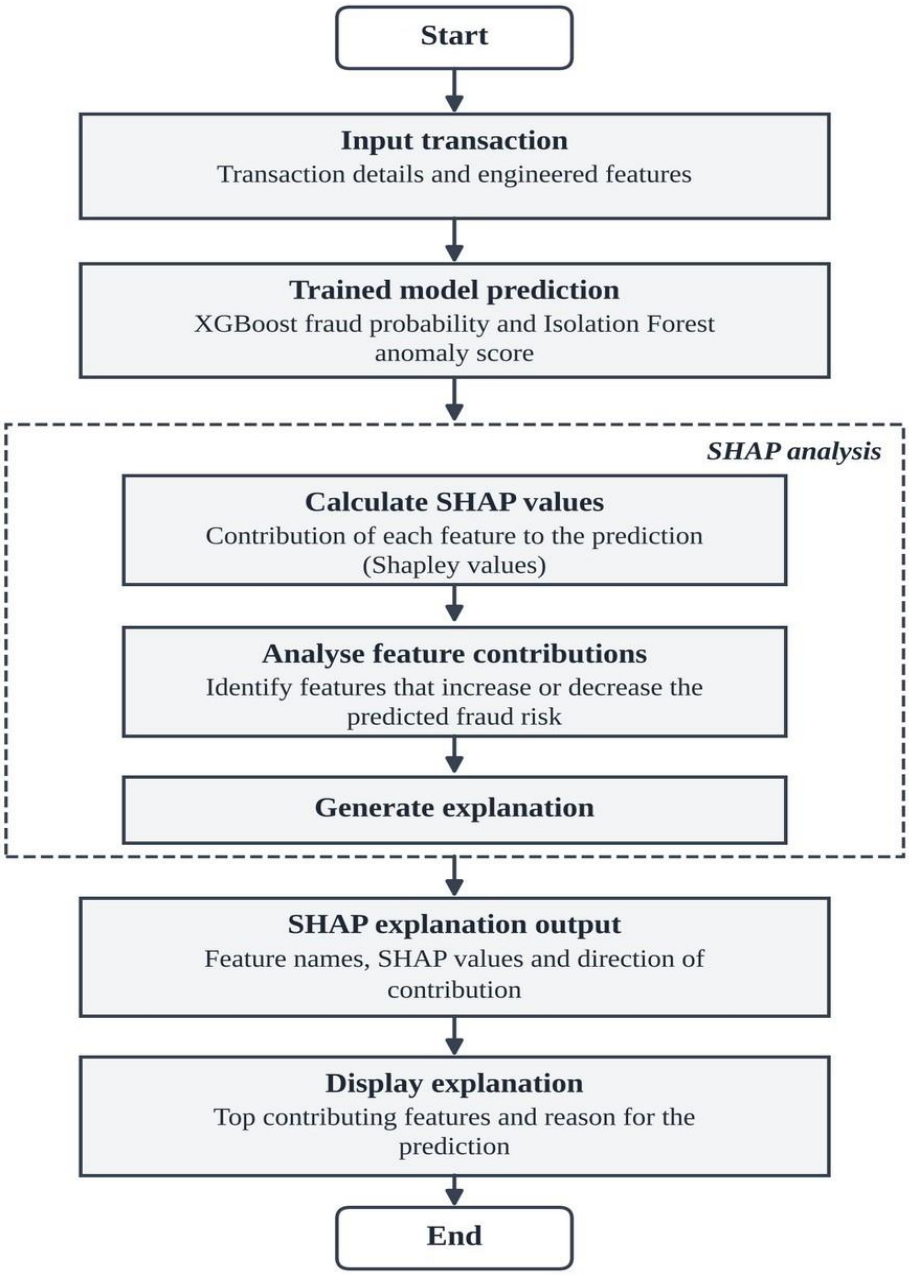

Fig. 6: Workflow of SHAP-based explanation generation for a model prediction.

5.10 Fraud Alert and Real-Time Notification System

When a transaction is classified as high risk, the system generates a fraud alert containing the transaction details, risk score, reason and timestamp, and stores it in the SQLite database. WebSockets then deliver a notification to the appropriate dashboards without requiring a manual refresh (Fig. 7). Alerts carry a status (Pending, Investigating or Resolved) that authorised users can review and track. Normal and medium-risk transactions are stored without generating an alert.

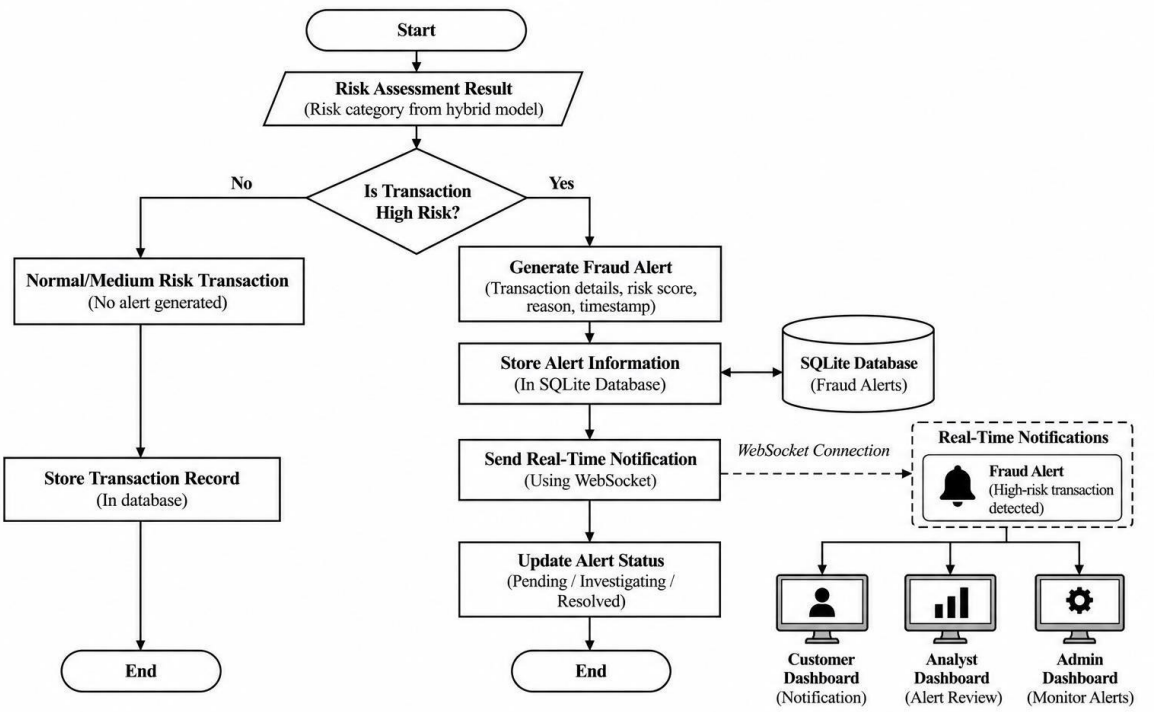

Fig. 7: Workflow of fraud alert generation and WebSocket-based real-time notification.

6. SYSTEM IMPLEMENTATION

6.1 Frontend Implementation

The frontend is a web application built with React.js and Vite, written in JavaScript and styled with CSS. The interface screens documented in this work are the login and registration pages (public registration creates Customer accounts), a customer portal for submitting transactions and viewing transaction history, and an administrator console with analytics and a WebSocket stream indicator. The integration status of the remaining views is summarised in Table 3.

6.2 Backend Implementation

The backend is written in Python with FastAPI. It provides the services for authentication, transaction submission, behavioural feature extraction, hybrid fraud prediction and risk scoring, SHAP explanation, alert generation and management, WebSocket notification, and PDF report generation. The backend was tested with Pytest, and 160 backend test cases passed. These are backend tests only; frontend and end-to-end testing are not covered in this paper.

6.3 Database Design

Application data is accessed through the SQLAlchemy object-relational mapper, and SQLite is the configured database. Four main entities are stored: **User** (credentials and role for authentication and access control), **Transaction** (submitted transaction details for monitoring

and analysis), **Fraud Alert** (suspicious transactions with risk details and alert status) and **User Behaviour Profile** (behavioural information used to identify unusual activity). These records are used by the fraud detection, monitoring, dashboard and reporting modules.

6.4 Authentication and Role-Based Access Control

Users log in with their credentials, which the backend verifies against the stored password hash; passwords are stored only in hashed form. After successful authentication the backend issues a JSON Web Token (JWT), and access to protected features is authorised according to the user's role (Fig. 8). Customers can access their own transactions, transaction status and related alerts. Analysts can monitor transactions, review fraud alerts, analyse risk details and generate reports. Administrators can manage users, monitor system activity and view all transactions and alerts.

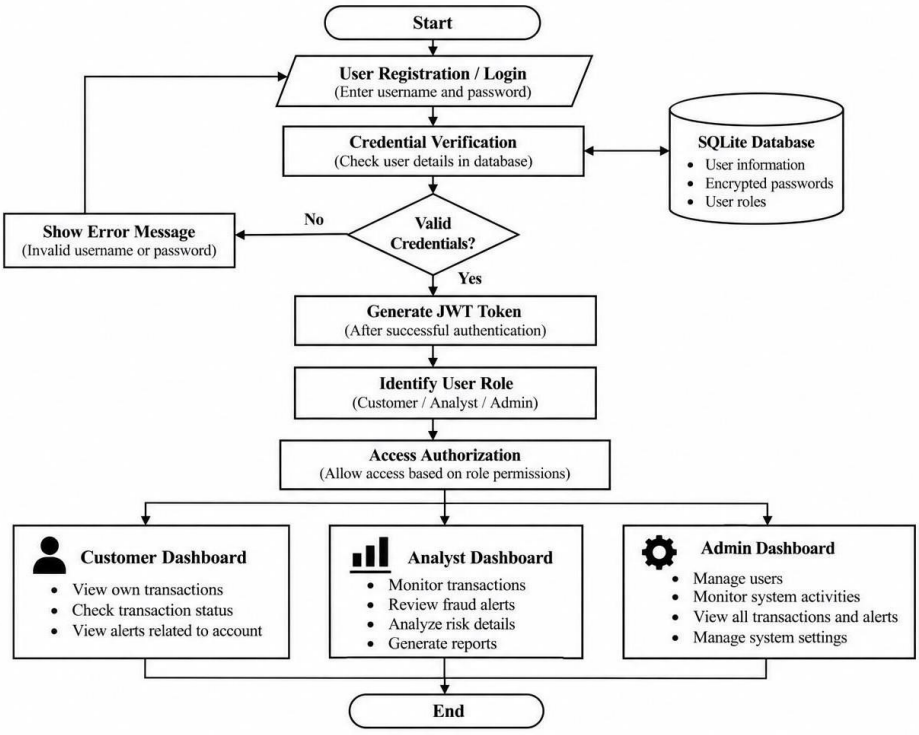

Fig. 8: Workflow of JWT authentication and role-based access for Customer, Analyst and Admin users.

6.5 Dashboard and Transaction Monitoring

Transactions are monitored through role-specific dashboards. The customer portal shows submitted transactions with their status and history. The administrator console shows the number of transactions, the fraud rate, the distribution of transactions by risk level, the average transaction amount and a trend of transaction volume and high-risk flags over time. Table 3 summarises the implementation status of the main functions as documented in the project

report.

Table 3: Implementation status of the main platform functions.

Function	Layer	Status
Authentication (JWT, password hashing) and role-based access	Backend	Implemented
Transaction processing and behavioural feature extraction	Backend	Implemented
Hybrid fraud prediction and Low/Medium/High risk scoring	Backend	Implemented
SHAP-based prediction explanations	Backend	Implemented
Fraud alert generation and WebSocket notification	Backend	Implemented
PDF audit report generation (ReportLab)	Backend	Implemented
Login, registration and customer portal (transaction submission, history)	Frontend	Implemented
Administrator analytics console with WebSocket stream indicator	Frontend	Implemented
Analyst dashboard, SHAP explanation view, alert management view and PDF report download	Frontend	Backend support implemented; interface integration in progress

6.6 PDF Audit Report Generation

Authorised users (Admin or Analyst) can request a PDF audit report. The backend retrieves transaction details, risk information, alert categories and alert status from the database, organises them into a structured report with headings and tables, and generates the PDF with the ReportLab library (Fig. 9). The report gives reviewers a record of transactions and fraud alerts for auditing.

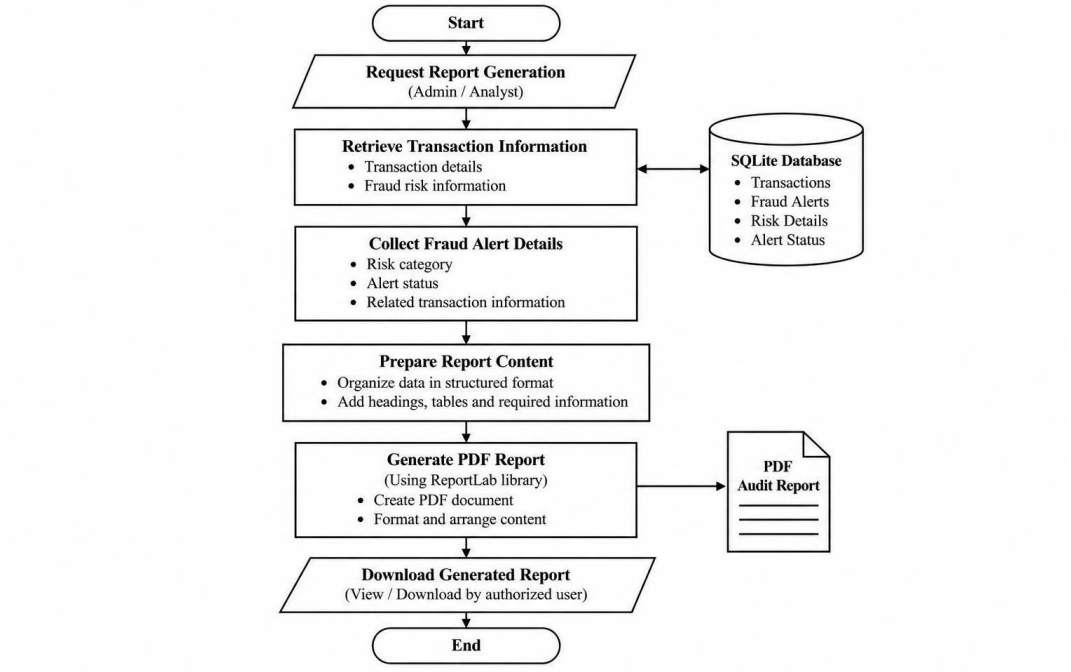

Fig. 9: Workflow of PDF audit report generation using ReportLab.

7. RESULTS AND PERFORMANCE EVALUATION

7.1 Dataset Description

The evaluation used the synthetic dataset described in Section 5.2 (Table 4). Only 5% of the records are fraudulent, which reflects the class imbalance typical of fraud detection and makes accuracy alone a poor measure of performance.

Table 4: Composition of the synthetic transaction dataset.

Class	Number of transactions	Share of dataset
Legitimate	4,750	95%
Fraudulent	250	5%
Total	5,000	100%

7.2 Experimental Results

Table 5 and Fig. 10 give the reported performance of the hybrid model on the synthetic dataset. The system performed transaction analysis, risk assessment and Low/Medium/High classification, and the backend supports alert generation, WebSocket notifications and PDF audit report generation. The metrics are aggregate values on the synthetic dataset.

Table 5: Reported performance of the hybrid model on the synthetic dataset.

Metric	Value	Primary focus
Precision	97.92%	Correctness of fraud predictions
Recall	94.00%	Coverage of actual fraud
F1-score	95.92%	Balance of precision and recall
ROC-AUC	99.98%	Overall class separation
PR-AUC	99.62%	Precision–recall trade-off

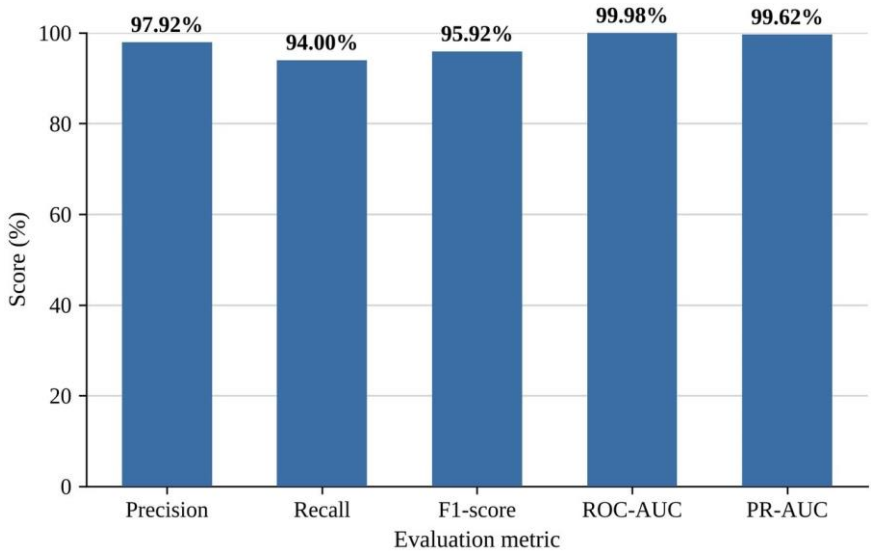

Fig. 10: Reported performance of the hybrid fraud detection model on the synthetic dataset (values as in Table 5).

7.3 Performance Metrics

Precision is the proportion of transactions predicted as fraudulent that are actually fraudulent. High precision means few legitimate transactions are wrongly flagged, which limits unnecessary alerts and customer inconvenience. **Recall** is the proportion of actual fraudulent transactions that the model identifies; it matters because every missed fraud is a potential financial loss. The **F1-score** is the harmonic mean of precision and recall and summarises the balance between them; the reported value of 95.92% is consistent with the reported precision and recall. **ROC-AUC** is the area under the curve relating true-positive and false-positive rates across all decision thresholds, and measures how well the model ranks fraudulent transactions above legitimate ones. **PR-AUC** is the area under the precision–recall curve; it focuses on the minority (fraud) class and is generally more informative than ROC-AUC when classes are imbalanced.

7.4 RESULT ANALYSIS AND DISCUSSION

The precision of 97.92% indicates that most transactions flagged as fraudulent were fraudulent, and the recall of 94.00% indicates that most of the fraudulent transactions in the evaluation data were detected. The F1-score of 95.92% reflects a good balance between the two. ROC-AUC (99.98%) and PR-AUC (99.62%) indicate strong separation between legitimate and fraudulent transactions on this dataset, and the high PR-AUC is meaningful because it accounts for the 5% fraud share.

These results must be interpreted with caution. First, the dataset is synthetic, and synthetic fraud patterns are often easier to separate than those in real transactions, so near-perfect AUC values should not be expected to carry over to real data. The results do not show that the platform would perform equally on real banking or UPI transactions. Second, only aggregate metrics are reported: the data partitioning and validation scheme, a confusion matrix and a comparison with single-model baselines (XGBoost alone or Isolation Forest alone) are not reported, so the benefit of the hybrid design over either model alone has not been quantified. Third, the platform has not been integrated with banking or payment systems or deployed in a production environment. Evaluation on real or publicly benchmarked data, with a documented split and baselines, is needed before any claim about practical effectiveness.

8. CONCLUSION

This paper presented an AI-based financial fraud detection platform that combines XGBoost classification and Isolation Forest anomaly detection over behavioural transaction features.

The platform assigns Low, Medium and High risk categories, explains predictions with SHAP, generates alerts delivered through WebSockets, and supports role-based access, transaction monitoring and PDF audit reporting through a React, FastAPI and SQLite implementation. The backend functions are implemented and 160 backend Pytest cases passed, while integration of some interface views is still in progress. On a synthetic dataset of 5,000 transactions the hybrid model achieved a precision of 97.92%, recall of 94.00%, F1-score of 95.92%, ROC-AUC of 99.98% and PR-AUC of 99.62%. Because these results come from synthetic data, they show that the approach works within that dataset but not how it would perform on real transactions, which remains to be evaluated.

9. FUTURE SCOPE

- **Real-world data.** Train and evaluate the models on larger real or benchmark transaction datasets, with a documented data split, baseline comparisons and a confusion matrix.
- **Frontend integration.** Complete the interface views for analysts, SHAP explanations, alert management and PDF report download on top of the existing backend services.
- **Live integration.** Integration with authorised banking and payment APIs could allow analysis of actual transactions as they occur; this has not been done in the present work.
- **Advanced models and behavioural features.** Additional ML and deep learning techniques, and richer user behaviour features, could be explored to detect more complex fraud patterns.
- **Alert management and deployment.** Advanced alert filtering and prioritisation would help analysts investigate suspicious transactions, cloud deployment could support larger transaction volumes, and a mobile application could give convenient access to monitoring and alerts.

REFERENCES

1. T. V. P. Monteiro, G. J. B. C. Castor, C. G. Castillo Correa, H. R. C. Arias, D. Z. Ñaupari Huatuco, and Y. P. Molina Rodriguez, “A Hybrid Machine Learning Framework for Electricity Fraud Detection: Integrating Isolation Forest and XGBoost for Real-World Utility Data,” *Energies*, vol. 18, no. 23, Art. no. 6249, 2025, doi: 10.3390/en18236249.
2. G. G. Battu, M. Garg, S. Dey, K. Kaushik, R. Kumawat, and M. Soni, “AI-Driven FinTech: Intelligent Risk Assessment and Fraud Detection in Digital Financial Ecosystems,” *Procedia Computer Science*, vol. 282, pp. 2667–2677, 2026, doi: 10.1016/j.procs.2026.06.036.

3. Z. Wang, Q. Shen, S. Bi, and C. Fu, “AI Empowers Data Mining Models for Financial Fraud Detection and Prevention Systems,” *Procedia Computer Science*, vol. 243, pp. 891–899, 2024, doi: 10.1016/j.procs.2024.09.107.
4. N. Uddin, “Role of AI in Preventing Financial Crime: A Comprehensive Analytical Review,” *Journal of Economic Criminology*, vol. 13, Art. no. 100200, 2026, doi: 10.1016/j.jeconc.2025.100200.
5. K. Elalkaoui, S. Moqqaddem, and M. Ait Kerroum, “A Comprehensive Study on Integration of Big Data and AI in Financial Decision-Making,” *Journal of ICT Standardization*, vol. 14, no. 2, pp. 199–222, 2026, doi: 10.13052/jicts2245-800X.1423.
6. Y. Lucas and J. Jurgovsky, “Credit Card Fraud Detection Using Machine Learning: A Survey,” arXiv preprint arXiv:2010.06479, 2020.
7. E. A. L. M. Btoush, X. Zhou, R. Gururajan, K. C. Chan, R. Genrich, and P. Sankaran, “A Systematic Review of Literature on Credit Card Cyber Fraud Detection Using Machine and Deep Learning,” *PeerJ Computer Science*, vol. 9, Art. no. e1278, 2023, doi: 10.7717/peerj-cs.1278.
8. M. Shanaa and S. Abdallah, “A Hybrid Anomaly Detection Framework Combining Supervised and Unsupervised Learning for Credit Card Fraud Detection,” *F1000Research*, vol. 14, Art. no. 664, 2025, doi: 10.12688/f1000research.166350.2.
9. N. Chapwanya and K. N. Gorejena, “Hybrid Unsupervised Machine Learning for Insurance Fraud Detection: PCA-XGBoost-LOF and Isolation Forest,” *Journal of Information Systems and Informatics*, vol. 7, no. 1, pp. 941–959, 2025, doi: 10.51519/journalisi.v7i1.958.
10. S. Dalal, B. Seth, M. Radulescu, C. Secara, and C. Tolea, “Predicting Fraud in Financial Payment Services through Optimized Hyper-Parameter-Tuned XGBoost Model,” *Mathematics*, vol. 10, no. 24, Art. no. 4679, 2022, doi: 10.3390/math10244679.