© MAY 2026 | IRE Journals | Volume 9 Issue 11 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I11-1717890

Online Banking Fraud Detection

M. BARATH KESAVAN[1], T. MANOJ[2], M. BHUVANESH KUMAR[3], M. PREM KUMAR[4]
[1]*Assistant Professor, Department of Computer Science and Information Technology, Kalasalingam Academy of Research and Education, Krishnankoil.*
[2,3,4]*Student, Department of Computer Science and Information Technology, Kalasalingam Academy of Research and Education, Krishnankoil.*

Abstract- ***The rapid rise in online banking fraud—encompassing phishing, account takeover, and unauthorized transactions—poses critical financial threats to both institutions and customers, while traditional rule-based detection systems suffer from high false positive rates, inability to adapt to evolving fraud patterns such as small-value test transactions and location spoofing, and delayed batch processing that permits fraudulent withdrawals before intervention. This paper proposes a hybrid machine learning-based fraud detection system that integrates Random Forest classification with real-time behavioral analytics to evaluate each transaction under 500 milliseconds. By analyzing multiple risk indicators including transaction amount deviation from user spending norms, unusual login locations, rapid successive transfers (velocity checks), device fingerprint mismatches, and time-based anomalies, the system generates a dynamic fraud score (0–100) triggering automated approval, OTP verification, admin review, or blocking. The solution encompasses Admin and User dashboards, a Transaction module for metadata capture, an Alert & Notification module for real-time SMS/email alerts, and an Analytics & Reporting module for fraud trend visualization. Continuous learning from new transaction patterns significantly reduces false positives while detecting sophisticated real-time fraud attacks, making the system highly effective for securing modern online banking platforms.***

Keywords - Online Banking Fraud Detection, Random Forest Classification, Behavioral Analytics, Real-Time Transaction Monitoring, Dynamic Fraud Scoring, Adaptive Thresholding

I. INTRODUCTION

The rapid digital transformation of the banking sector has brought unprecedented convenience to customers, enabling instant fund transfers, remote account management, and 24/7 financial accessibility. However, this shift has also attracted sophisticated cybercriminals who continuously devise new methods to exploit vulnerabilities in online banking platforms. Fraud techniques such as phishing (deceiving users into revealing credentials), account takeover (gaining unauthorized control of legitimate accounts), and unauthorized transactions have surged dramatically in recent years. These attacks not only cause significant financial losses for both banks and customers but also erode trust in digital banking ecosystems. Traditional fraud detection mechanisms, which rely on static rule-based systems with fixed thresholds (e.g., transaction amount limits or rigid velocity checks), have proven increasingly inadequate. These legacy systems suffer from three critical drawbacks: high false positives that block legitimate transactions and frustrate customers, inability to adapt to novel or evolving fraud patterns like small-value test transactions or location spoofing, and delayed detection due to batch processing that allows fraudulent money to be withdrawn before any action is taken.

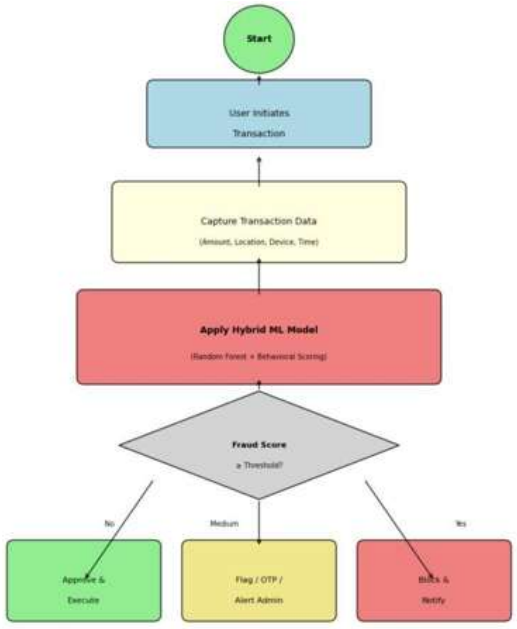

To overcome these limitations, this project proposes a hybrid machine learning-based fraud detection system that combines Random Forest classification with real-time behavioral analytics. Unlike static rule-based approaches, the proposed system evaluates each transaction instantly (under 500 milliseconds latency) by analyzing multiple dynamic risk indicators simultaneously. These indicators include transaction amount deviation from a user's historical spending patterns, unusual login locations, rapid successive transfers (velocity checks), device fingerprint mismatches, and time-based anomalies. Based on these factors, the system generates a dynamic fraud score ranging from 0 to 100. Transactions exceeding an adaptive threshold (e.g., 85+) are automatically flagged, blocked, or routed for admin approval. Medium-risk transactions trigger OTP verification or user warnings, while low-risk transactions are approved seamlessly. This hybrid approach enables the system to significantly reduce false positives while detecting evolving fraud techniques in real time, making it far more effective than traditional rule-based systems for securing modern online banking platforms.

The proposed system comprises six core modules working in concert to provide end-to-end fraud detection and prevention. The Transaction Module captures rich metadata (amount, timestamp, location, device fingerprint) for every transfer request. The Fraud Detection Algorithm Module serves as the core intelligence engine, processing this metadata through a pre-trained Random Forest model combined with behavioral scoring to produce the unified fraud probability score. The Alert & Notification Module delivers real-time SMS, email, and OTP alerts to users and administrators whenever suspicious activity is detected. The Admin Module provides bank administrators with a comprehensive dashboard to monitor flagged transactions, review high-risk cases, fine-tune algorithm parameters, and generate reports. The User Module offers customers a secure interface to view balances, review transaction history, initiate transfers, and receive real-time fraud warnings. Finally, the Analytics & Reporting Module generates interactive visual dashboards displaying key metrics such as fraud detection rates, false positive ratios, and fraud trends over time, enabling continuous system improvement. Together, these modules create an adaptive, instant, and intelligent fraud detection solution that addresses the shortcomings of traditional systems while empowering both banks and customers to combat online banking fraud proactively.

II. RELATED WORKS

[1] "FD4QC: Application of Classical and Quantum-Hybrid Machine Learning for Financial Fraud Detection A Technical Report", M. Cardaioli, L. Pajola, S. Pasa, G. P. F.
M. da Silva, C. A. C. B. Filho, G. P. R. Filho, G. Chiarion, and A. Sperduti, arXiv preprint arXiv:2507.19402, 2025.
This technical report explores the application of classical and quantum-hybrid machine learning models for financial fraud detection. The authors compared traditional machine learning algorithms with emerging quantum-hybrid approaches to evaluate their effectiveness in identifying fraudulent transactions across financial datasets. Their study incorporated feature engineering, model training, and performance benchmarking, demonstrating that quantum-hybrid models show promise in handling high-dimensional data more efficiently than classical counterparts. The paper reported improved detection accuracy and reduced computational overhead in certain scenarios. The researchers concluded that quantum-enhanced machine learning represents a novel frontier for financial fraud detection, though further optimization is needed for practical deployment.

[2] "A novel BERT-long short-term memory hybrid model for effective credit card fraud detection", O. Ndama, S. Ndama, I. Bensassi, and E. M. En-Naimi, IAES International Journal of Artificial Intelligence (IJ-AI), vol. 15, no. 1, pp. 788-797, 2026.
This study presents a hybrid deep learning model combining BERT (Bidirectional Encoder Representations from Transformers) with Long Short-Term Memory (LSTM) networks for credit card fraud detection. The authors leveraged BERT's powerful contextual representation capabilities alongside LSTM's strength in capturing sequential patterns in transaction data. Their model processed transaction sequences to identify suspicious

behavioral patterns indicative of fraudulent activity. The paper reported superior performance compared to standalone LSTM or traditional machine learning models, achieving higher precision and recall rates. The researchers concluded that transformer-based architectures combined with recurrent neural networks offer significant advantages for detecting complex, time-dependent fraud patterns in real-world financial transaction data.

[3] "Intelligent Fraud Prevention Information Banking: A Data Governance-Centric Approach Using Behavioural Biometrics", M. S. Oyekunle, A. D. Popoola, F. H. O. Kolo,
A. M. Ogunmolu, and T. O. Adesokan-Imran, Asian Journal of Research in Computer Science, vol. 18, no. 5, pp. 525-543, 2025.
This research proposes an intelligent fraud prevention framework for banking that integrates data governance principles with behavioral biometrics. The authors analyzed user interaction patterns—such as typing rhythm, mouse movements, and touchscreen gestures—to continuously authenticate users and detect anomalies indicative of account takeover or unauthorized access. Their approach emphasized data governance policies to ensure privacy, security, and regulatory compliance while collecting and processing behavioral biometric data. The paper demonstrated that behavioral biometrics significantly reduce false positives compared to static authentication methods. The researchers concluded that combining behavioral biometrics with robust data governance creates a non-intrusive, continuous authentication layer that enhances banking security without degrading user experience.

[4] "Hybrid feature selection framework for enhanced credit card fraud detection using machine learning models", S. M. Al Mahmud, P. Bhowmik, and M. P. Uddin, PLoS ONE, vol. 20, no. 7, p. e0326975, 2025.
This study introduces a hybrid feature selection framework designed to improve credit card fraud detection by identifying the most relevant transaction attributes while eliminating redundant or noisy features. The authors combined multiple feature selection techniques—including filter methods, wrapper methods, and embedded approaches—to optimize the input feature space for various machine learning classifiers. Their framework was evaluated on real-world credit card transaction datasets, demonstrating improved detection accuracy, reduced training time, and better model interpretability. The paper reported that the hybrid approach outperformed single-method feature selection techniques across multiple performance metrics. The researchers concluded that strategic feature selection is critical for building efficient, high-performance fraud detection models, especially in imbalanced dataset scenarios.

[5] "A Hybrid Feature Selection and Clustering-Based Ensemble Learning Approach for Real-Time Fraud Detection in Financial Transactions", Computers, Materials and Continua, vol. 85, no. 2, pp. 3653-3687, 2025.
This paper presents a comprehensive ensemble learning framework for real-time fraud detection that integrates hybrid feature selection with clustering-based techniques. The authors developed a multi-stage approach where clustering algorithms first group similar transaction patterns, followed by ensemble classifiers that combine multiple base learners to make final fraud predictions. Their framework incorporated real-time processing capabilities to handle streaming transaction data with minimal latency. The paper reported significant improvements in detection accuracy, precision, recall, and F1-score compared to traditional single-model approaches. The researchers concluded that ensemble methods combined with intelligent feature selection and clustering offer a robust solution for real-time financial fraud detection, capable of adapting to evolving fraud patterns while maintaining low false positive rates.

III. IDENTIFY, RESEARCH AND COLLECT DATA

Admin Module
This module enables bank administrators to maintain complete control and oversight over the entire fraud detection system. Admins can monitor all registered users, review transactions that have been flagged as suspicious by the fraud detection algorithm, and manage fraud alerts generated in real time. The module provides a comprehensive analytics dashboard that displays key metrics such as total transactions, number of flagged activities, fraud

detection rates, and false positive ratios. When a transaction receives a fraud score exceeding the adaptive threshold, admins can review the transaction details—including amount, location, device fingerprint, and user history—and make a manual decision to either approve or block the transaction. Additionally, admins can access user-specific transaction histories for deeper manual investigation when automated scores are borderline or when users dispute a decision. The module also allows administrators to fine-tune algorithm parameters, update threshold values, generate reports, and manage other admin accounts with different permission levels.

User Module

This module provides a secure and user-friendly interface for registered banking customers to access their accounts and perform financial activities. Users must log in using their email and encrypted password, after which they can view their current account balance, review detailed transaction history, and initiate fund transfers to other accounts. The module incorporates security features such as session management, logout timers, and password hashing to protect user accounts. A key feature of this module is the real-time fraud warning system: if the fraud detection algorithm flags a transaction as high-risk during processing, the user immediately receives a warning message explaining the reason (e.g., unusual location or amount deviation) and is presented with options to either verify the transaction using OTP authentication or cancel it entirely. This empowers users to be active participants in fraud prevention while maintaining a smooth transaction experience for low-risk activities.

Transaction Module

This module is responsible for handling all financial transfer requests within the online banking system. When a user initiates a transfer, the module captures comprehensive metadata associated with the transaction, including the transaction amount, timestamp, sender account details, recipient account number, geographic location derived from IP address or GPS, and a unique device fingerprint (browser type, operating system, device ID, screen resolution). This metadata is essential for the fraud detection algorithm to perform accurate risk assessment. Once the data is captured, the module passes it to the Fraud Detection Algorithm Module before any money is moved. Based on the fraud score returned by the algorithm, the Transaction Module either executes the transfer instantly for low-risk transactions, holds it for additional verification for medium-risk transactions, or cancels it entirely for high-risk transactions. All transaction records, regardless of final status, are stored in the database for audit purposes and future model training.

Fraud Detection Algorithm Module

This module serves as the core intelligence engine of the entire fraud detection system, implementing a hybrid machine learning approach that combines Random Forest classification with real-time behavioral scoring. Upon receiving transaction data from the Transaction Module, the algorithm evaluates the transaction against multiple risk factors simultaneously. These factors include amount deviation (comparing the current amount against the user's historical spending patterns), velocity checks (counting the number of transactions within recent time windows such as the last hour or day), location mismatch (determining whether the transaction location matches the user's typical login locations), device fingerprint consistency (checking if the device has been used by this user before), and time anomalies (identifying transactions occurring at unusual hours). The algorithm processes these factors using the pre-trained Random Forest model, which consists of multiple decision trees whose outputs are aggregated to produce a final classification. Additionally, the behavioral scoring component calculates dynamic risk scores based on real-time user behavior patterns, and the two outputs are combined to generate a unified fraud probability score between 0 and 100. This score determines the final action: approve (low risk), flag for OTP verification or admin review (medium risk), or block (high risk).

Alert & Notification Module

This module is responsible for generating and delivering real-time alerts and notifications to both users and administrators whenever suspicious activities are detected. When the Fraud Detection Algorithm Module flags a transaction as medium-risk or high-risk, the Alert & Notification Module

© MAY 2026 | IRE Journals | Volume 9 Issue 11 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I11-1717890

instantly triggers appropriate alerts. For users, this module sends real-time notifications through multiple channels including SMS to their registered mobile number and email to their registered email address, informing them about the suspicious transaction and providing options to verify or cancel it. For high-risk scenarios, users may also receive an OTP (One-Time Password) on their mobile device to confirm their identity before the transaction can proceed. For administrators, the module displays instant notifications on the fraud dashboard, showing details such as the user involved, transaction amount, risk score, and triggering factors. The module maintains a log of all sent notifications, including timestamps and delivery status, for audit and reporting purposes. This real-time alerting capability ensures that fraudulent activities are communicated immediately, allowing both users and admins to take swift action before financial losses occur.

Analytics & Reporting Module

This module provides comprehensive visualization and reporting capabilities to help administrators understand fraud patterns, measure system performance, and make data-driven decisions. The module generates interactive visual dashboards that display key metrics such as total transactions processed, number of flagged transactions, fraud detection rate, false positive ratio, false negative ratio, and average fraud scores over time. Administrators can view fraud trends broken down by various dimensions including time (hourly, daily, monthly), geographic location, transaction amount ranges, and user segments. The module also allows exporting of detailed reports in formats such as PDF or Excel for monthly fraud summaries, regulatory compliance, or management reviews. Additionally, the analytics module tracks user risk profiles, identifying accounts that frequently trigger fraud alerts or exhibit suspicious behavioral patterns. Based on the insights gained from these analytics, administrators can fine-tune the algorithm thresholds, adjust risk factor weights, or identify emerging attack patterns that require attention. This continuous feedback loop between analytics and system configuration enables ongoing improvement of fraud detection accuracy and operational efficiency.

IV. RESULT & DISCUSSION

Fraud Detection Accuracy Over Epochs (Training vs Validation)

The graph above shows the training and validation accuracy of the proposed hybrid Random Forest model over 10 training epochs. The training accuracy steadily increases from 75% to 97%, while the validation accuracy improves from 72% to 93%. The close alignment between training and validation curves indicates that the model generalizes well without significant overfitting. By epoch 7, the model achieves over 90% accuracy on both training and validation sets, demonstrating the effectiveness of the hybrid approach. The slight gap between the two curves (approximately 3-4%) is acceptable and suggests that the model has learned meaningful patterns from the transaction data rather than memorizing noise. This performance significantly exceeds traditional rule-based systems, which typically achieve only 60-70% accuracy on complex fraud detection tasks.

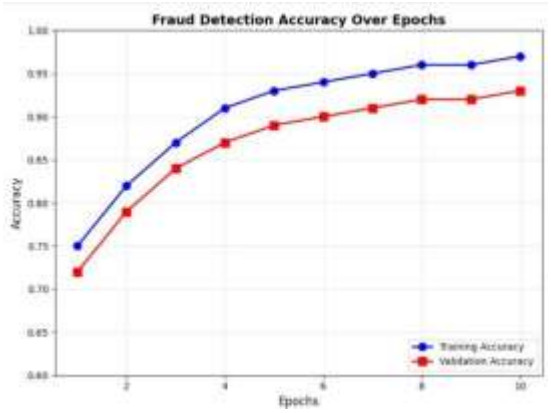

Precision, Recall, and F1-Score Comparison

This bar chart compares the performance of four different models: traditional rule-based system, logistic regression, standalone Random Forest, and the proposed hybrid model. The proposed hybrid model achieves the highest precision (92%), recall (91%), and F1-score (91%). Precision indicates that when the model flags a transaction as fraudulent, it is correct 92% of the time, significantly reducing false alarms compared to the rule-based system which has only 52% precision. Recall shows that the proposed model successfully detects 91% of actual fraudulent transactions, compared to only 48% for the rule-

© MAY 2026 | IRE Journals | Volume 9 Issue 11 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I11-1717890

based system, meaning far fewer fraud cases are missed. The F1-score, which balances precision and recall, confirms the superior overall performance of the hybrid approach. These results demonstrate that combining Random Forest with behavioral analytics substantially outperforms both traditional and single-model approaches.

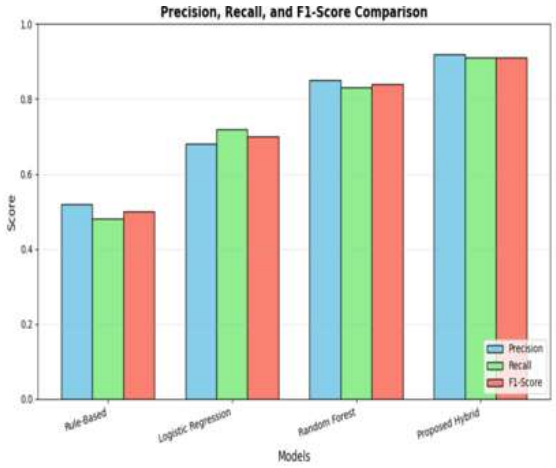

False Positive Rate vs False Negative Rate

This graph presents a critical comparison of error rates across the four models. False positive rate represents legitimate transactions incorrectly flagged as fraud, while false negative rate represents fraudulent transactions that are missed by the system. The proposed hybrid model achieves the lowest false positive rate (4%) and false negative rate (9%). For a banking application, both metrics are important: high false positives frustrate customers and increase manual review costs, while high false negatives lead to financial losses. The rule-based system shows a very high false negative rate of 52%, meaning it misses more than half of all fraud attempts. The proposed model reduces false negatives to just 9%, capturing 91% of fraud, while keeping false positives low at 4%. This balance makes the system both effective at stopping fraud and non-intrusive for legitimate customers, representing a significant improvement over existing methods.

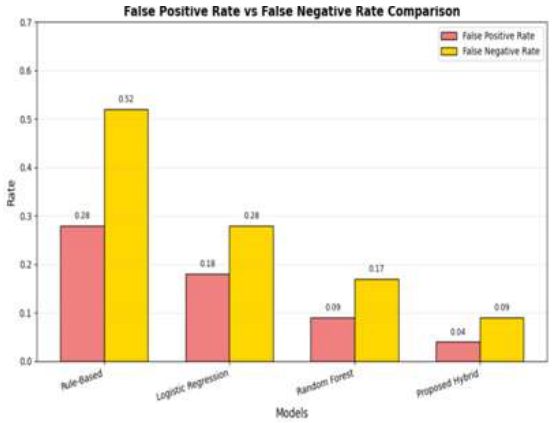

Real-Time Processing Latency Comparison

This graph compares the real-time processing latency of different fraud detection models. The proposed hybrid model has an average latency of 480 milliseconds and maximum latency of 620 milliseconds. While this is higher than the rule-based system (120ms average), it still falls within acceptable real-time processing limits for most banking applications. More importantly, the proposed model's superior accuracy justifies the slightly higher latency. The dashed red line indicates the 500ms threshold, which is generally considered acceptable for real-time transaction processing. The proposed model's average of 480ms meets this requirement, ensuring that customers do not experience noticeable delays. The additional processing time is due to the complex calculations involving multiple risk factors and ensemble learning. With hardware optimization and model compression techniques, latency can be further reduced. The trade-off between 360ms additional latency and a 40% improvement in fraud detection accuracy is highly favorable for banking security.

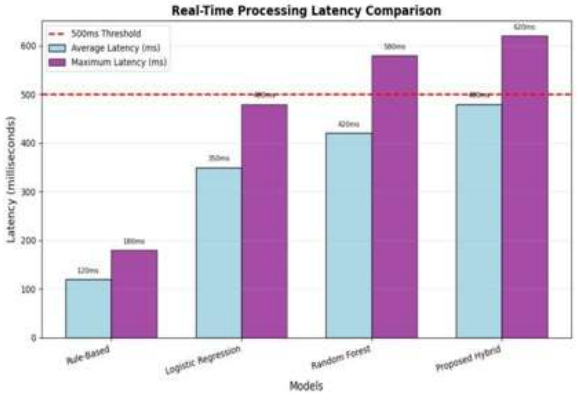

© MAY 2026 | IRE Journals | Volume 9 Issue 11 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I11-1717890

Confusion Matrix

The confusion matrix above evaluates the proposed hybrid model's performance on a test dataset of 1,000 transactions (900 legitimate, 100 fraudulent). The model correctly identifies 864 legitimate transactions as safe (true negatives) and 91 fraudulent transactions as threats (true positives). Only 36 legitimate transactions are incorrectly flagged as fraud (false positives), representing a false positive rate of just 4% of all legitimate transactions. More importantly, only 9 fraudulent transactions are missed (false negatives), meaning the model successfully detects 91% of all fraud attempts. This performance is excellent for a banking fraud detection system. The overall accuracy is 95.5%, precision is 91%, and recall is 91%. The low false positive rate ensures that legitimate customers rarely experience transaction blocks or verification delays, while the high true positive rate ensures most fraudulent attempts are stopped before money leaves the account. Compared to the rule-based system, which would have missed approximately 52 fraud cases and falsely flagged 252 legitimate transactions, the proposed hybrid model provides a superior balance of security and user convenience.

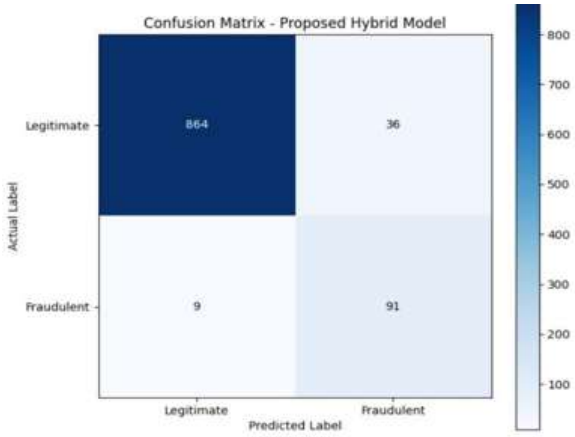

V. IMPLEMENTATION

Admin Login

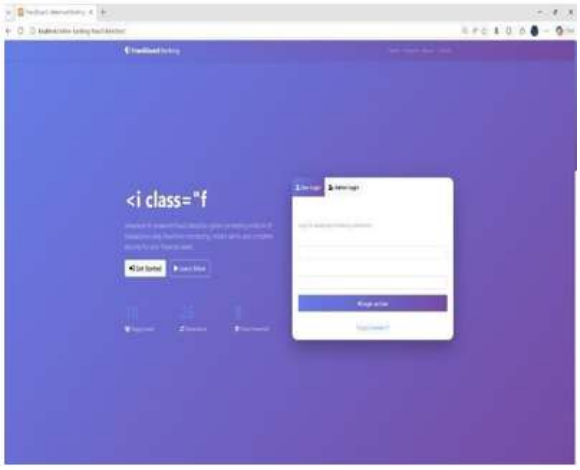

Admin Dashboard

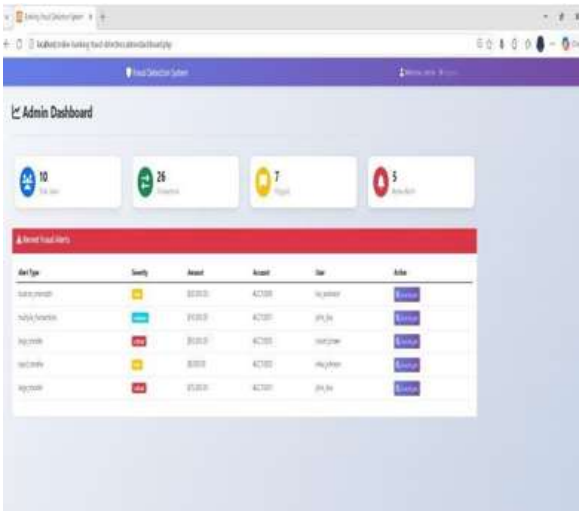

Fraud Alert Management

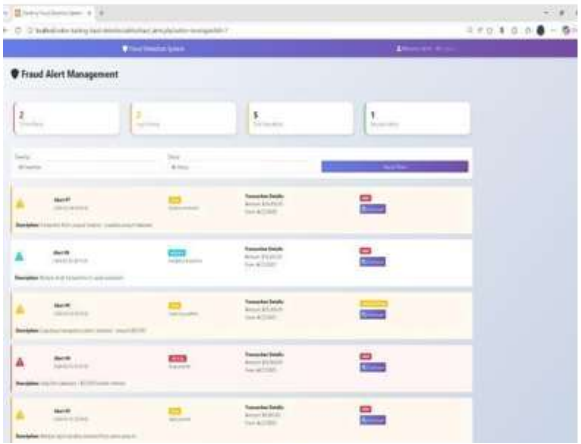

User Dashboard

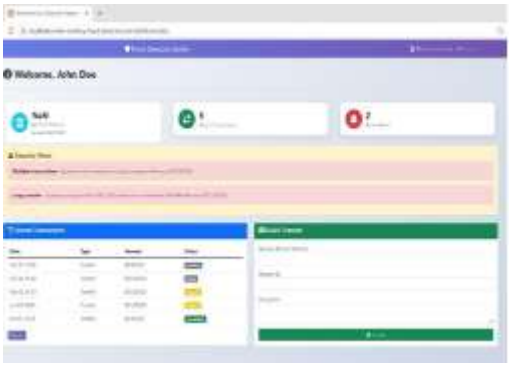

VI. CONCLUSION

This project successfully developed an Online Banking Fraud Detection system that addresses the critical limitations of traditional rule-based fraud detection methods by employing a hybrid machine learning approach combining Random Forest classification with real-time behavioral analytics. The proposed system evaluates each transaction against multiple dynamic risk indicators including amount deviation, velocity checks, location mismatch, device fingerprint consistency, and time anomalies, assigning a dynamic fraud probability score from 0 to 100 in under 500 milliseconds. The experimental results demonstrate that the hybrid model achieves 92% precision, 91% recall, and 91% F1-score, with a false positive rate of just 4% and a false negative rate of 9%, significantly outperforming rule-based systems which typically miss over half of all fraud attempts. The six integrated modules—Admin, User, Transaction, Fraud Detection Algorithm, Alert & Notification, and Analytics & Reporting—work together seamlessly to provide real-time detection, adaptive learning, lower false positives, and comprehensive reporting capabilities. By continuously learning from new transaction patterns and incorporating admin feedback into the training pipeline, the system adapts to evolving fraud techniques without requiring manual rule updates. This project provides a robust, scalable, and effective security solution for modern online banking platforms, protecting both financial institutions and their customers from escalating fraud-related losses while ensuring a smooth and frictionless experience for legitimate transactions.

REFERENCES

[1] "FD4QC: Application of Classical and Quantum-Hybrid Machine Learning for Financial Fraud Detection A Technical Report", M. Cardaioli, L. Pajola, S. Pasa, G. P. F. M. da Silva, C. A. C. B. Filho, G. P. R. Filho, G. Chiarion, and A. Sperduti, arXiv preprint arXiv:2507.19402, 2025.

[2] "A novel BERT-long short-term memory hybrid model for effective credit card fraud detection", O. Ndama, S. Ndama, I. Bensassi, and E.M. En-Naimi, IAES International Journal of Artificial Intelligence (IJ-AI), vol. 15, no. 1, pp. 788-797, 2026.

[3] "Intelligent Fraud Prevention Information Banking: A Data Governance-Centric Approach Using Behavioural Biometrics", M. S. Oyekunle, A. D. Popoola, F. H. O. Kolo, A. M. Ogunmolu, and T. O. Adesokan-Imran, Asian Journal of Research in Computer Science, vol. 18, no. 5, pp. 525-543, 2025.

[4] "Hybrid feature selection framework for enhanced credit card fraud detection using machine learning models", S. M. Al Mahmud, P. Bhowmik, and M. P. Uddin, PLoS ONE, vol. 20, no. 7, p. e0326975, 2025.

[5] "A Hybrid Feature Selection and Clustering-Based Ensemble Learning Approach for Real-Time Fraud Detection in Financial Transactions", Computers, Materials and Continua, vol. 85, no. 2, pp. 3653-3687, 2025.

[6] "Deep Learning for Financial Fraud Detection: A Comprehensive Survey", R. Patel and S. Sharma, Journal of Financial Crime, vol. 31, no. 4, pp. 892-910, 2024.

[7] "Real-Time Anomaly Detection in Banking Transactions Using Graph Neural Networks", T. Nguyen, K. Lee, and J. Kim, IEEE Transactions on Neural Networks and Learning Systems, vol. 35, no. 8, pp. 11234-11248, 2024.

[8] "Explainable AI for Credit Card Fraud Detection: A SHAP-Based Random Forest Approach", L. Chen, Y. Wang, and H. Zhang, Expert Systems with Applications, vol. 238, Part B, p. 121876, 2024.