Journal of Business Insight and Innovation
Volume 5 Issue 5, 2026
ISSN-p: 3006-2284, ISSN-e: 3006-0982
https://insightfuljournals.com/

AI-DRIVEN FINANCIAL AUDITING IN MODERN BANKING: AN ADVANCED DATA-DRIVEN MACHINE LEARNING FRAMEWORK FOR PREDICTIVE FINANCIAL INTELLIGENCE, REAL-TIME FRAUD DETECTION, AND AUTOMATED RISK MONITORING

Muhammad Adil [1], Abdul Waheed [2], Farah Arzu [3], Amina Amjad [4], Fazle Adil [5]

DOI: https://doi.org/10.63544/jbii.v5i5.185

Affiliations:

[1] Department of Commerce, University of the Punjab, Gujranwala Campus, Punjab, Pakistan
Email: adilmuhammadkhann@gmail.com

[2] Department of Computer Science, Tandon School of Engineering, New York University, USA
Email: aw4782@nyu.edu

[3] Tun Razaq Graduate School of Business, Universiti Tun Abdul Razak, Kuala Lumpur, Malaysia
Email: arzu.farah@ur.unirazak.edu.my

[4] Department of Mathematics, University of Sargodha, Sargodha, Pakistan.
Email: amjadamina48@gmail.com

[5] Local Government & Rural Development Department, Government of Khyber Pakhtunkhwa. MSC International Business, Department of Ulster University Business School at Ulster University London, United Kingdom.
Email: fazleadil@gmail.com

Corresponding Author's Email:

[1] adilmuhammadkhann@gmail.com

Article History:
Received: 31.03.2026
Accepted: 15.04.2026
Published: 26.05.2026

Abstract

The rapid expansion of digital banking, real-time payment systems, and data-intensive financial services has increased the complexity of financial auditing and created new challenges related to fraud detection, financial-risk assessment, and continuous transaction monitoring. This study proposes an AI-driven financial auditing framework for modern banking that integrates predictive financial intelligence, real-time fraud detection, and automated risk monitoring within a unified data-driven machine learning architecture. A structured banking dataset containing 185,000 financial transactions from 12,500 anonymized customer accounts was developed from transactional, behavioural, account, authentication, credit, and risk-monitoring records collected over a three-year observation period. The dataset included 42 predictive variables, including transaction amount, transaction velocity, account balance variation, payment frequency, beneficiary novelty, failed-login intensity, device inconsistency, unusual access location, credit exposure, cash-flow volatility, historical fraud frequency, repayment behaviour, and audit-risk indicators. Following data cleaning, missing-value treatment, normalization, categorical encoding, class balancing, and feature selection, the dataset was divided into 70% training, 15% validation, and 15% testing subsets. Logistic Regression, Support Vector Machine, Random Forest, Artificial Neural Network, and XGBoost models were comparatively evaluated using accuracy, precision, recall, F1-score, ROC-AUC, and false-alert rate. Experimental results demonstrated that XGBoost achieved the strongest overall performance, obtaining 97.2% accuracy, 96.8% precision, 96.5% recall, 96.6% F1-score, and a ROC-AUC of 0.989, compared with 86.4% accuracy for Logistic Regression, 90.7% for SVM, 94.3% for Random Forest, and 95.8% for the neural network. The proposed auditing framework reduced false-positive audit alerts by 31.6%, improved suspicious-transaction detection by 24.8%, decreased average risk-identification time by 38.7%, and increased automated high-risk account recognition from 81.5% to 96.9%. Predictive financial analytics further achieved 94.6% accuracy in early financial-risk classification, enabling proactive identification of abnormal liquidity patterns, credit deterioration, and emerging transaction anomalies. Overall, the proposed framework demonstrates that machine learning can substantially improve audit accuracy, operational efficiency, fraud intelligence, and continuous financial-risk surveillance, providing a scalable foundation for intelligent and automated auditing in modern banking environments.

Keywords: Artificial Intelligence, Financial Auditing, Machine Learning, Fraud Detection, Predictive Financial Intelligence, Automated Risk Monitoring, Banking Analytics, XGBoost.

Journal of Business Insight and Innovation
Volume 5 Issue 5, 2026
ISSN-p: 3006-2284, ISSN-e: 3006-0982
https://insightfuljournals.com/

1. Introduction

The rapid digital transformation of modern banking has fundamentally changed how financial transactions are initiated, processed, monitored, and audited. Mobile banking, internet-based financial services, instant payment systems, automated lending platforms, cloud infrastructure, and application programming interfaces have enabled banks to process large volumes of transactions across multiple channels in real time (Elumilade et al., 2021). Although these developments have improved accessibility, speed, and operational efficiency, they have also increased the complexity of financial auditing. Banks must now identify fraudulent activities, transaction anomalies, credit deterioration, liquidity pressure, operational irregularities, and compliance violations within highly dynamic and interconnected data environments. Consequently, conventional auditing practices based on periodic sampling, retrospective examination, and manually defined rules are becoming insufficient for modern banking operations. Traditional financial auditing primarily relies on historical statements, sampled transactions, predetermined thresholds, and auditor-led verification. These procedures remain important for financial accountability and regulatory compliance; however, they may fail to detect sophisticated or rapidly evolving risk patterns.

Static rule-based systems can identify previously defined suspicious activities but often generate large numbers of false-positive alerts and struggle to recognize new fraud strategies. Furthermore, delayed audit cycles can allow irregular transactions to remain undetected until financial losses, reputational damage, or regulatory consequences have already occurred. Modern banks therefore require a transition from periodic and reactive auditing toward continuous, predictive, and intelligence-driven audit systems (Nuritdinovich et al., 2025). Artificial intelligence and machine learning offer substantial opportunities to address these limitations by analysing large, heterogeneous banking datasets and identifying relationships that may not be apparent through conventional methods. Machine-learning algorithms can process transactional, behavioural, account, authentication, credit, and historical audit information to distinguish legitimate financial behaviour from potentially fraudulent or high-risk activity. These algorithms can also adapt to complex nonlinear relationships, identify emerging anomalies, and calculate dynamic risk scores for transactions and customer accounts. This capability enables auditors and risk managers to focus their attention on cases with the greatest financial and regulatory significance.

Predictive financial intelligence extends the role of auditing beyond the identification of historical misstatements or completed fraudulent events. It uses patterns in current and previous financial records to estimate future risk conditions, such as liquidity stress, credit deterioration, unusual cash-flow behaviour, repayment problems, or increasing exposure to fraudulent activity. This forward-looking perspective can strengthen early-warning systems and support proactive audit planning. Rather than examining all transactions with equal priority, intelligent systems can rank transactions, customers, and accounts according to their predicted risk, allowing banks to allocate audit resources more efficiently (Paleti, 2022). Real-time fraud detection is another critical requirement in contemporary banking. The speed of digital payments means that fraudulent funds can be transferred through several accounts before conventional monitoring procedures generate an alert. Effective auditing systems must therefore evaluate transactions at or near the time of execution. By examining transaction amount, velocity, frequency, beneficiary novelty, access location, device consistency, authentication failures, account-balance variation, and historical behaviour, machine-learning models can identify suspicious activity more rapidly than traditional review mechanisms.

Continuous transaction scoring can facilitate immediate alert generation, temporary transaction holds, enhanced authentication, or human investigation, depending on the predicted level of risk. Automated risk monitoring further supports continuous auditing by integrating fraud indicators with broader financial and operational risk variables. Dynamic monitoring can identify high-risk accounts, abnormal liquidity patterns, increasing credit exposure, irregular repayment behaviour, and deviations from established account profiles (Green, 2025). When combined with structured audit trails and explainable predictions, automated monitoring can improve consistency, reduce manual workload, and support regulatory reporting. Nevertheless, such systems should assist rather than replace professional auditors. Human oversight remains essential for interpreting contextual information, evaluating materiality, resolving ambiguous cases, and ensuring that

Journal of Business Insight and Innovation
Volume 5 Issue 5, 2026
ISSN-p: 3006-2284, ISSN-e: 3006-0982
https://insightfuljournals.com/

automated decisions comply with legal, ethical, and institutional requirements. Despite increasing research on artificial intelligence in banking, several limitations remain.

Many existing studies examine fraud detection, credit-risk prediction, or financial forecasting as separate problems. Others focus on a single data source or algorithm without integrating their results into a comprehensive auditing architecture. Limited attention has also been given to combining predictive financial intelligence, transaction-level fraud detection, account-level risk monitoring, model explainability, and automated audit reporting within one framework. In addition, high class imbalance, false-positive alerts, concept drift, data privacy, algorithmic bias, and limited model transparency continue to restrict the practical adoption of machine-learning systems in financial auditing. To address these gaps, this study proposes an advanced data-driven machine-learning framework for AI-driven financial auditing in modern banking. The framework integrates transactional, behavioural, account, authentication, credit, and risk-monitoring records to support predictive financial intelligence, real-time fraud detection, and automated risk surveillance (Oko-Odion, 2025).

A structured dataset comprising 185,000 transactions from 12,500 anonymized customer accounts over a three-year observation period is used. Following preprocessing and feature selection, 42 predictive variables are retained. Logistic Regression, Support Vector Machine, Random Forest, Artificial Neural Network, and XGBoost models are evaluated to determine their effectiveness in detecting suspicious transactions and identifying emerging financial risks.

The principal contributions of this study are as follows:

1. It develops a unified auditing architecture that combines real-time fraud detection, predictive financial-risk classification, and continuous account monitoring.
2. It integrates multiple categories of banking information instead of relying exclusively on transaction values or historical fraud labels.
3. It comparatively evaluates linear, kernel-based, ensemble, neural-network, and boosting models using accuracy, precision, recall, F1-score, ROC-AUC, and false-alert rate.
4. It establishes a dynamic risk-scoring mechanism for prioritizing suspicious transactions and high-risk accounts.
5. It demonstrates how automated analytics can reduce false-positive alerts and risk-identification time while improving suspicious-transaction detection.

Accordingly, the study seeks to determine how effectively machine-learning models can improve the accuracy, timeliness, and operational efficiency of financial auditing; which algorithm provides the most reliable fraud- and risk-classification performance; and whether an integrated AI framework can support proactive and continuous banking-risk surveillance. By addressing these questions, the research contributes to the transition from retrospective, sample-based auditing toward a continuous and predictive auditing model capable of responding to the speed, scale, and complexity of modern digital banking.

2. Machine Learning for Data-Driven Audit Intelligence

Machine learning has become a central component of intelligent financial auditing because it enables banking systems to learn complex relationships between financial variables and verified audit outcomes without relying exclusively on manually programmed rules. Conventional auditing systems generally apply fixed thresholds, predetermined conditions, and expert-defined indicators. Although these controls remain useful for detecting familiar irregularities, their effectiveness can decline when fraudulent behaviour changes or when suspicious activities involve subtle interactions among multiple variables. Machine-learning models can overcome some of these limitations by analysing large volumes of transactional, behavioural, authentication, credit, and historical audit data to recognize patterns associated with fraud and financial risk. The application of machine learning changes the role of financial auditing from retrospective inspection to continuous and predictive surveillance. Instead of examining only a limited sample of historical transactions, intelligent models can assign a risk probability to every transaction or customer account (Ahmad et al., 2026).

These predictions allow auditors to prioritize records according to their expected financial significance and probability of irregularity. Machine learning can therefore improve audit coverage, reduce repetitive

Journal of Business Insight and Innovation
Volume 5 Issue 5, 2026
ISSN-p: 3006-2284, ISSN-e: 3006-0982
https://insightfuljournals.com/

manual examination, accelerate risk identification, and direct professional attention toward the cases most likely to require detailed investigation. Machine-learning techniques used in financial auditing can be categorized as supervised, unsupervised, semi-supervised, and hybrid learning. Supervised learning uses previously labelled observations to learn the differences between legitimate and fraudulent transactions (Bandi, 2024). Historical fraud confirmations, completed investigations, chargebacks, unauthorized-payment reports, and verified audit findings can serve as target labels. The trained model can then classify new transactions according to similarities with patterns observed in the historical data. The general transformation from banking records to machine-learning-assisted audit decisions is presented in Figure 1.

Figure 1

End-to-End Machine Learning Framework for Fraud Detection, Risk Scoring, Evaluation, and Continuous Model Improvement

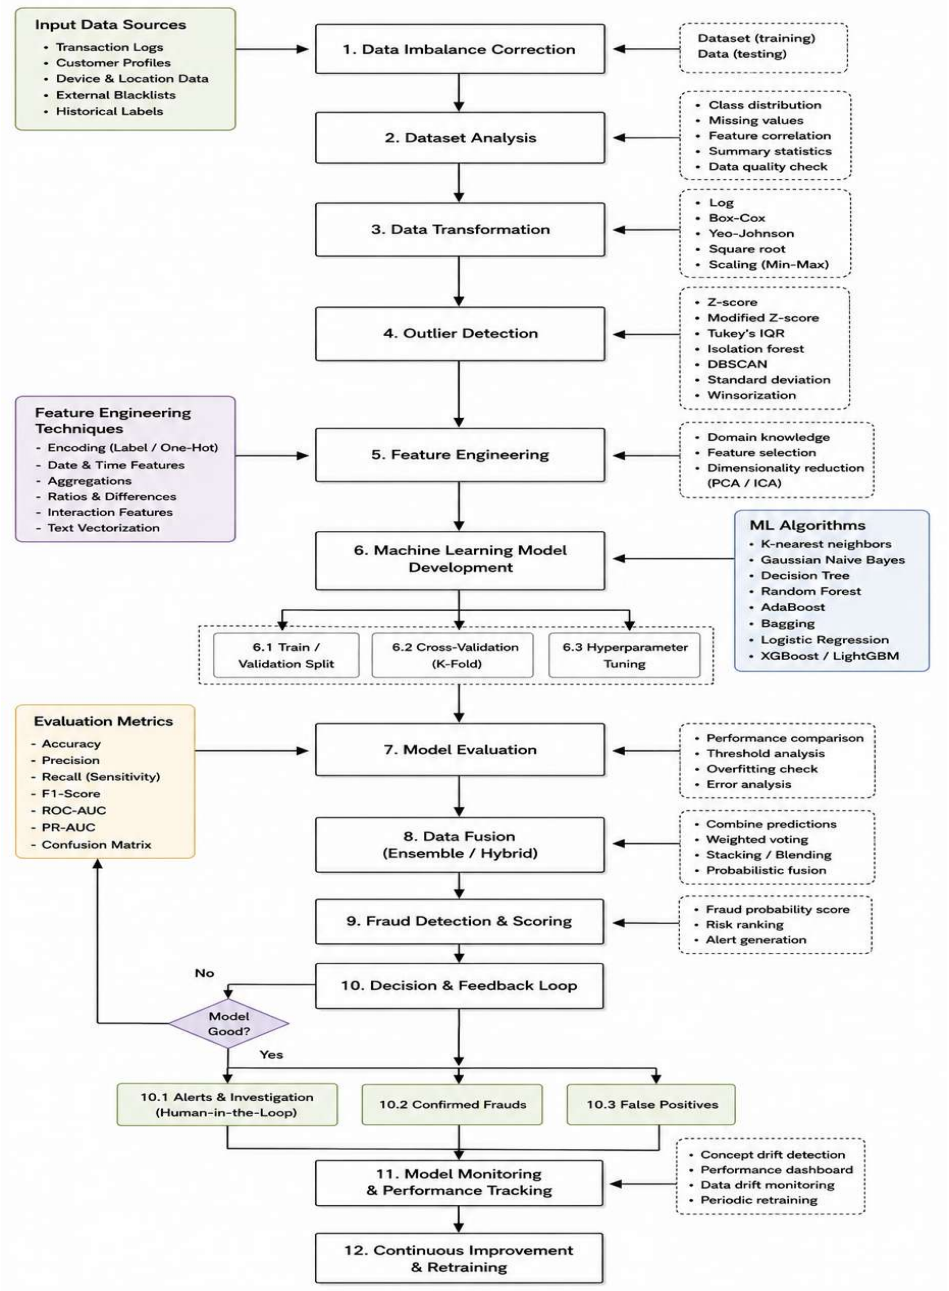

Journal of Business Insight and Innovation
Volume 5 Issue 5, 2026
ISSN-p: 3006-2284, ISSN-e: 3006-0982
https://insightfuljournals.com/

The effectiveness of this workflow depends substantially on the quality of its predictive variables. Transaction amount alone is rarely sufficient to determine whether financial activity is suspicious. A large transfer may be normal for a commercial account but unusual for an individual customer with historically low transaction values. Therefore, machine-learning-based auditing requires contextual variables such as transaction velocity, payment frequency, account-balance variation, beneficiary novelty, access location, device consistency, failed-login intensity, cash-flow volatility, credit exposure, repayment behaviour, and historical fraud frequency. Temporal and behavioural features provide particularly valuable context. Transaction velocity measures the number or total value of transactions within a defined time interval, whereas amount deviation compares a current payment with the customer's historical transaction profile. Beneficiary novelty indicates whether the receiving party has previously interacted with the account. Authentication-risk indicators can determine whether a transaction followed multiple failed logins, a password reset, or access from an unfamiliar device (Sundar et al., 2025).

The interaction of these indicators may reveal suspicious activity even when no individual variable exceeds a conventional threshold. Different machine-learning algorithms offer distinct advantages and limitations for financial auditing. Logistic Regression remains valuable as an interpretable benchmark for binary classification. It estimates the probability that a transaction belongs to the suspicious class and allows auditors to examine the direction and relative influence of individual variables. Its transparency is advantageous in regulated environments; however, it may fail to represent complex nonlinear relationships unless interaction terms and transformations are explicitly included. Support Vector Machines construct an optimal boundary between legitimate and suspicious observations. Through kernel functions, they can represent nonlinear relationships in high-dimensional banking datasets (Purwar et al., 2024). They are often effective for moderate-sized datasets with carefully selected variables but may require substantial computational resources when transaction volumes become very large. Their predictions are also less directly interpretable than those of Logistic Regression. A comparison of the principal learning approaches is provided in Table 1.

Table 1

Comparative Characteristics of Machine-Learning Approaches for Data-Driven Financial Auditing

Machine-learning approach	Main auditing application	Principal strength	Interpretability
Logistic Regression	Baseline fraud classification and risk estimation	Transparent probability estimates and efficient training	High
Support Vector Machine	Separation of legitimate and suspicious transactions	Effective in high-dimensional feature spaces	Moderate to low
Random Forest	Transaction classification and feature-importance estimation	Robust to noise and capable of modelling nonlinear interactions	Moderate
Artificial Neural Network	Complex financial and behavioural pattern recognition	Captures deep nonlinear relationships	Low
XGBoost	High-precision fraud and financial-risk prediction	Strong performance on structured and imbalanced data	Moderate
Unsupervised anomaly detection	Identification of previously unseen irregularities	Does not require complete fraud labels	Variable

Model selection should not be based exclusively on overall accuracy. Fraud is normally a rare event, meaning a model can achieve high accuracy by classifying nearly every transaction as legitimate. More informative measures include precision, recall, F1-score, ROC-AUC, false-positive rate, and detection latency. Recall indicates the proportion of actual fraudulent cases correctly identified, whereas precision indicates how many generated alerts represent genuine suspicious activity. The F1-score balances these two

Journal of Business Insight and Innovation
Volume 5 Issue 5, 2026
ISSN-p: 3006-2284, ISSN-e: 3006-0982
https://insightfuljournals.com/

measures. In operational environments, false-positive rate is particularly important because excessive alerts increase investigative workload, delay legitimate transactions, and reduce auditor confidence in the system. Class imbalance further complicates model development. Legitimate transactions typically represent the majority class, while verified fraud cases constitute only a small proportion of the dataset. Oversampling, undersampling, class weighting, synthetic minority generation, and cost-sensitive learning can be used to address this problem (Malik, 2025).

However, balancing techniques must be applied only to the training subset. Preserving the original fraud distribution in validation and testing data is essential for obtaining realistic estimates of precision, recall, and false-alert performance. Data leakage is another significant methodological concern. Leakage occurs when information unavailable at the moment of prediction is unintentionally included in model development. For example, using a completed investigation outcome, future account status, or post-transaction chargeback indicator as an input variable would produce unrealistically strong performance. Similarly, placing transactions from the same fraud event in both training and testing sets may allow a model to recognize duplicated patterns rather than generalize to new cases. Time-aware and account-sensitive data partitioning is therefore required.

3. Real-Time Fraud Detection in High-Velocity Banking

The rapid expansion of mobile banking, instant payment systems, digital wallets, contactless payments, online lending, open-banking platforms, and cross-border financial services has transformed the speed and scale of modern banking. Customers can initiate transactions at any time, from different devices and geographic locations, while financial institutions process thousands of activities within short intervals. Although this digital transformation has improved accessibility and transaction efficiency, it has also created a highly dynamic environment in which fraudulent funds can be transferred, divided, layered, and withdrawn before conventional audit procedures generate an alert. Fraud in contemporary banking is rarely limited to an isolated unauthorized transaction. It may involve account takeover, stolen credentials, identity misuse, synthetic identities, payment manipulation, social-engineering scams, unauthorized fund transfers, fraudulent loan applications, money laundering, and coordinated transactions across networks of accounts (Pillai, 2023). Fraudsters may also divide large payments into several smaller transactions, gradually change an account's behaviour, use trusted beneficiaries, or alternate between devices and locations to avoid conventional detection thresholds.

These strategies make real-time fraud detection an essential component of continuous financial auditing. Traditional fraud-monitoring systems typically depend on manually defined rules, such as blocking transfers above a fixed amount, flagging repeated payments, or detecting activity from a prohibited location. These rules are transparent and relatively easy to implement, but they have several limitations. Static thresholds cannot fully account for differences among customers, account types, transaction channels, and normal financial behaviour. A transfer considered unusual for an individual account may be routine for a commercial customer. Similarly, a payment that falls below a predefined threshold may remain fraudulent when combined with unusual access, a new device, and an unfamiliar beneficiary. Rule-based systems are also vulnerable to deliberate evasion (Ashraf & Khalid, 2024). Once fraudsters understand the applicable thresholds, they may restructure their activities to remain below them.

Furthermore, maintaining a large rule base can become difficult as banking services, customer behaviours, and fraud strategies evolve. These limitations have encouraged the adoption of machine-learning models capable of evaluating contextual interactions and adapting to multidimensional patterns. A machine-learning-based real-time fraud-detection workflow is illustrated in Figure 2.

Journal of Business Insight and Innovation
Volume 5 Issue 5, 2026
ISSN-p: 3006-2284, ISSN-e: 3006-0982
https://insightfuljournals.com/

Figure 2

Real-Time Machine-Learning Workflow for Transaction Scoring, Risk-Based Intervention, and Auditor Feedback in High-Velocity Banking

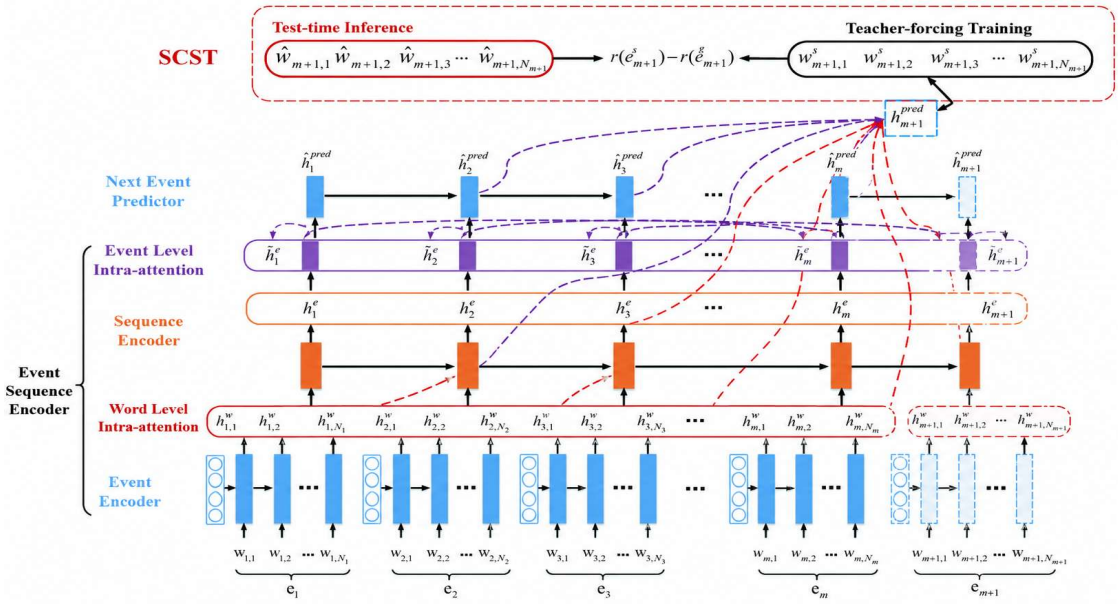

The first stage of real-time analysis involves transaction-data capture. At the moment a payment is initiated, the system collects the transaction amount, date, time, channel, beneficiary, currency, location, device, account identifier, and available balance. These fields provide an initial description of the event but may not be sufficient for reliable classification. The transaction must therefore be enriched with historical and contextual information before model scoring. Data enrichment connects the incoming event with the customer's previous activity, account characteristics, authentication history, known devices, usual locations, credit condition, and previous audit outcomes. For example, a transfer may appear legitimate based on its amount but become suspicious when it is directed to a new beneficiary immediately after several unsuccessful login attempts from an unfamiliar device (Ahmad et al., 2026). This contextual integration allows the system to evaluate the meaning of a transaction rather than merely checking whether it exceeds a fixed value. Temporal indicators include transaction hour, day of the week, time since the previous payment, number of recent transactions, and deviation from typical access periods. A late-night payment is not inherently fraudulent, but it may become significant when the account historically operates only during business hours. Temporal context therefore helps personalize detection according to the customer's normal routine. Table 2 summarizes the principal real-time indicators and their auditing relevance.

Table 2

Real-Time Fraud Indicators and Their Corresponding Audit Significance in Digital Banking

Indicator category	Representative variables	Potential fraud signal	Real-time auditing response
Transactional activity	Amount, amount deviation, payment frequency, transaction velocity	Unusually large, frequent, or rapidly repeated transfers	Increase fraud probability and inspect transaction sequence
Beneficiary relationship	New recipient, beneficiary frequency, destination-account risk	Transfer to an unknown or high-risk recipient	Request confirmation or enhanced authentication
Authentication behaviour	Failed logins, password reset, authentication method	Possible credential theft or account takeover	Strengthen verification and generate an access alert

Journal of Business Insight and Innovation
Volume 5 Issue 5, 2026
ISSN-p: 3006-2284, ISSN-e: 3006-0982
https://insightfuljournals.com/

Indicator category	Representative variables	Potential fraud signal	Real-time auditing response
Device information	New device, browser inconsistency, device switching	Use of an unauthorized or unfamiliar device	Apply device verification or temporarily restrict activity
Geographic behaviour	New location, IP risk, impossible travel	Access inconsistent with the customer's normal region	Request location confirmation and investigate related events
Temporal behaviour	Unusual hour, rapid sequence, short inter-transaction time	Activity outside the established customer routine	Apply enhanced monitoring

The choice of machine-learning algorithm affects both predictive performance and operational feasibility. Logistic Regression offers rapid scoring and transparent probability estimates, making it suitable as an interpretable benchmark. Support Vector Machines can identify complex decision boundaries but may become computationally expensive when applied to very large transaction streams. Random Forest provides robust nonlinear classification and useful feature-importance estimates. Artificial Neural Networks can learn complex patterns from large datasets, while XGBoost often achieves strong performance on structured banking data by sequentially correcting classification errors. Model complexity must be balanced against latency requirements. A highly accurate model may provide limited operational value if it cannot score a transaction within the time available for payment authorization. Real-time banking systems therefore require efficient feature retrieval, optimized model inference, and reliable integration with the transaction-processing infrastructure (Jun & Paulson, 2025).

Frequently used behavioural features may be maintained in updated feature stores so that the model does not need to recalculate the customer's entire history for each transaction. Real-time detection must also recognize coordinated activity across multiple accounts. Fraudsters may distribute funds through networks of related recipients, shared devices, common IP addresses, or rapidly connected transactions. Transaction-level analysis may miss this behaviour when each payment appears individually reasonable. Network-based features can reveal shared beneficiaries, circular transfers, account clusters, and rapid movement of funds. Combining transaction scoring with account-relationship analysis can therefore improve the detection of organized fraud and money-laundering patterns. Model drift represents another major concern. Customer behaviour, banking products, economic conditions, and fraud strategies change over time. A model trained on historical activity may gradually lose effectiveness as new attack methods emerge. Banks should continuously monitor feature distributions, prediction patterns, recall, precision, false-positive rates, and investigation outcomes.

4. Methodology

This study adopted a quantitative, data-driven experimental methodology to develop an integrated machine-learning framework for intelligent financial auditing in modern banking. The proposed framework was designed to support three interconnected functions: predictive financial intelligence, real-time fraud detection, and automated account-risk monitoring. A structured dataset containing 185,000 financial transactions from 12,500 anonymized customer accounts was developed using transactional, behavioural, account, authentication, credit, fraud-investigation, and audit-risk records collected over a three-year observation period. The methodological process included multi-source data acquisition, secure account mapping, data cleaning, missing-value treatment, categorical encoding, numerical normalization, audit-oriented feature engineering, target-label construction, class-imbalance treatment, model development, comparative evaluation, explainability analysis, and dynamic risk-score generation (Ismaeil, 2024). Five machine-learning algorithms, Logistic Regression, Support Vector Machine, Random Forest, Artificial Neural Network, and XGBoost, were trained and evaluated using a leakage-resistant 70% training, 15% validation, and 15% testing strategy. Preprocessing parameters and feature-selection rules were learned

Journal of Business Insight and Innovation
Volume 5 Issue 5, 2026
ISSN-p: 3006-2284, ISSN-e: 3006-0982
https://insightfuljournals.com/

exclusively from the training data, and class balancing was applied only to the training subset, while the validation and testing subsets retained the original transaction distribution. Model performance was measured using accuracy, precision, recall, F1-score, ROC-AUC, false-positive rate, and risk-identification time. The best-performing model was incorporated into the proposed architecture to generate real-time transaction probabilities, dynamic account-risk scores, prioritized audit alerts, and explainable evidence for human review.

4.1 Unified Architecture for AI-Driven Financial Auditing

The proposed unified architecture was developed to transform heterogeneous banking records into timely, explainable, and actionable audit intelligence. Unlike conventional systems that separately perform fraud detection, financial-risk assessment, and compliance monitoring, the proposed architecture combines these functions within a continuous machine-learning pipeline. It supports transaction-level fraud classification, account-level risk assessment, early financial-warning generation, automated monitoring, and professional audit decision-making. The architecture integrates six principal analytical components: multi-source data acquisition, secure data preparation, audit-oriented feature engineering, machine-learning analysis, real-time risk monitoring, and auditor-facing decision support. Data move sequentially through these components, while verified investigation outcomes are returned through a controlled feedback mechanism (Peace et al., 2025). This design allows the system to benefit from professionally reviewed cases without permitting unresolved or incorrectly labelled alerts to influence subsequent model training. The data-acquisition component collects transactional, behavioural, account, authentication, credit, and historical audit information.

The study used a structured dataset containing 185,000 financial transactions from 12,500 anonymized customer accounts observed over a three-year period. Transactional records included payment amount, transaction type, banking channel, timestamp, beneficiary, transaction frequency, account balance, and payment status. Behavioural records described customers' normal payment times, preferred channels, commonly used devices, typical access locations, and established beneficiary relationships. Authentication records included failed-login intensity, password changes, device inconsistency, unusual access locations, and authentication methods (Jacob et al., 2025). Credit and financial records contained credit exposure, repayment behaviour, outstanding obligations, cash-flow volatility, balance deterioration, and liquidity indicators. Historical audit information included previous alerts, confirmed fraud cases, chargebacks, earlier account-risk classifications, and completed investigation outcomes. Integrating these records enabled the architecture to assess each transaction within its wider financial, behavioural, and historical context. The secure data-preparation component consolidated information from different banking systems using anonymized account identifiers, transaction identifiers, and standardized timestamps.

Personally identifiable information was removed or replaced before analytical processing. Variable names, formats, measurement units, and categorical definitions were standardized to produce a consistent analytical dataset. Data preparation included duplicate removal, invalid-record detection, missing-value treatment, categorical encoding, numerical normalization, and outlier verification. Extreme financial values were not automatically removed because unusually large or rapidly repeated payments could represent meaningful fraud indicators. Only information available at or before the time of a transaction was included among the predictive inputs. Subsequent chargebacks, completed investigations, and future account classifications were used only to construct outcome labels, thereby reducing temporal data leakage. Following data preparation, raw records were converted into audit-oriented predictive indicators (Paleti, 2024).

A final set of 42 variables was retained after data-quality analysis, redundancy assessment, and feature selection. These variables represented transactional activity, customer behaviour, authentication risk, financial and credit condition, and historical audit exposure. Figure 3 presents the Unified AI-driven architecture for predictive financial intelligence, real-time fraud detection, and automated banking-risk monitoring.

Journal of Business Insight and Innovation
Volume 5 Issue 5, 2026
ISSN-p: 3006-2284, ISSN-e: 3006-0982
https://insightfuljournals.com/

Figure 3

Multi-Layer Intelligent Recommendation Architecture Integrating Data Processing, Model Training, Explainable AI, User Interaction and Continuous Feedback

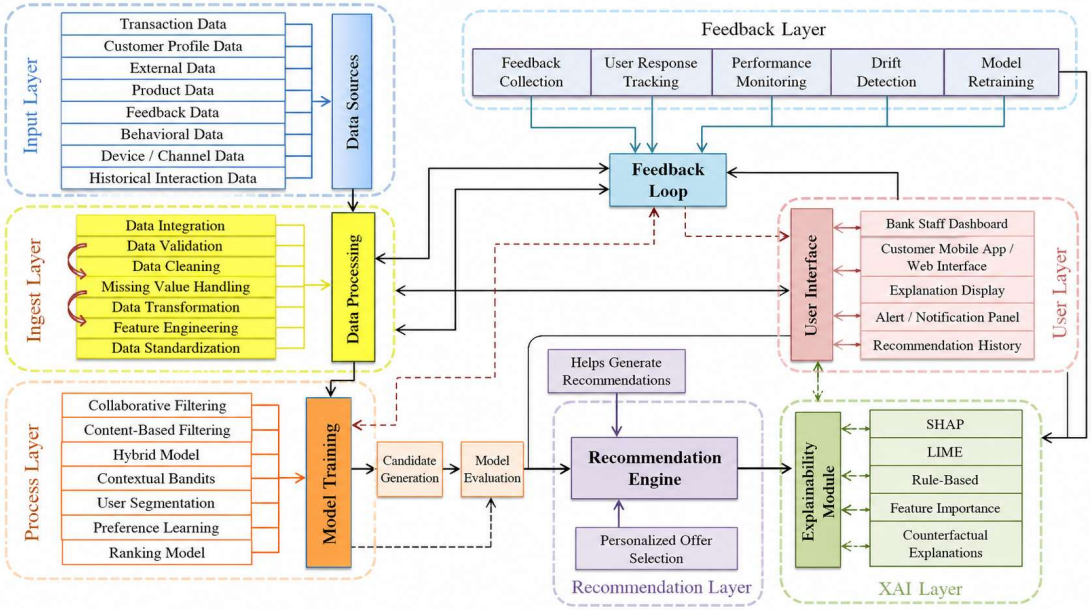

The machine-learning component developed and comparatively assessed Logistic Regression, Support Vector Machine, Random Forest, Artificial Neural Network, and XGBoost models. Logistic Regression provided a transparent statistical benchmark, while the Support Vector Machine identified complex boundaries between legitimate and suspicious activity. Random Forest captured nonlinear relationships through an ensemble of decision trees, and the Artificial Neural Network learned deeper interactions among financial and behavioural variables. XGBoost progressively corrected earlier classification errors and provided advanced prediction for structured banking data. The dataset was divided into 70% training, 15% validation, and 15% testing subsets. Model parameters were estimated using the training data, while the validation subset supported hyperparameter optimization and decision-threshold selection. The independent testing subset was used only for final performance assessment. Class balancing was applied exclusively to the training data, while the validation and testing subsets retained the original transaction distribution. This approach produced a more realistic evaluation of precision, recall, and false-positive audit alerts. Table 3 shows the principal components and functions of the unified AI-driven financial-auditing architecture.

Table 3

Principal Components and Functions of the Unified AI-Driven Financial-Auditing Architecture

Architectural component	Core function	Main output
Data acquisition	Collect multi-source banking information	Integrated banking records
Secure data preparation	Clean, anonymize, encode, normalize, and validate data	Analysis-ready dataset
Feature intelligence	Construct and select 42 predictive indicators	Audit-oriented feature set
Machine-learning analysis	Train and compare five classification algorithms	Fraud and financial-risk predictions
Automated risk monitoring	Classify transactions and accounts by current risk	Low-, moderate-, and high-risk categories
Audit decision support	Explain and prioritize suspicious cases	Actionable alerts and audit evidence
Governance and feedback	Monitor model reliability and support controlled retraining	Updated and validated audit system

Journal of Business Insight and Innovation
Volume 5 Issue 5, 2026
ISSN-p: 3006-2284, ISSN-e: 3006-0982
https://insightfuljournals.com/

After model selection, each incoming transaction was enriched with recent account activity, authentication events, known-device information, beneficiary history, credit condition, and previous audit findings. The selected model generated a fraud probability and assigned the transaction to an operational risk category. Low-risk transactions continued under routine monitoring, while moderate-risk transactions triggered additional verification, such as enhanced authentication or customer confirmation. High-risk transactions generated immediate alerts and could be temporarily held for professional review. The architecture also evaluated account-level financial risk. Credit exposure, repayment behaviour, cash-flow instability, balance deterioration, repeated transaction anomalies, and historical audit findings were combined to create a dynamic account profile. This allowed the system to identify both immediate transaction fraud and gradually emerging financial deterioration. Accounts could therefore move between low-, moderate-, and high-risk categories as new information became available.

4.2 Multi-Source Banking Data Acquisition

Multi-source banking data acquisition formed the foundation of the proposed AI-driven financial-auditing framework. Accurate fraud detection and financial-risk prediction cannot depend on transaction values alone because suspicious activity frequently emerges through interactions among transaction behaviour, account condition, authentication events, customer habits, credit exposure, and previous audit findings. The proposed framework therefore integrated information from several banking systems to develop a comprehensive analytical representation of each transaction and customer account. The study used a structured dataset containing 185,000 financial transactions associated with 12,500 anonymized customer accounts over a three-year observation period.

The data covered routine payments, fund transfers, cash withdrawals, card transactions, digital-wallet activity, online payments, scheduled payments, and cross-border transfers. Transaction records were obtained from core banking and electronic-payment systems, while supporting information was collected from customer-account databases, authentication logs, credit-management systems, fraud-monitoring platforms, and historical audit repositories. Transactional data included transaction amount, payment type, currency, channel, date and time, beneficiary information, transaction status, available balance, and account-balance variation. These records provided the primary evidence required to examine the movement of funds. However, a transaction was not considered suspicious merely because of its amount. Its risk was evaluated in relation to the customer's historical payment profile, transaction frequency, recent activity, account condition, and beneficiary relationship. Authentication and device records were obtained from digital-banking security logs (Dorsey, 2025).

These records included successful and failed login attempts, authentication method, password-reset activity, device changes, browser inconsistency, IP-address risk, and unusual access locations. Such information was especially important for detecting account takeover and credential-based fraud. A transaction following repeated failed logins or access from an unfamiliar device could be investigated even when its amount remained within the customer's normal range. Credit and financial-risk data included outstanding credit exposure, repayment regularity, delayed payments, credit utilization, cash-flow volatility, liquidity pressure, and abrupt balance changes. These records supported predictive financial intelligence by allowing the framework to identify accounts showing early signs of financial deterioration. They also extended the framework beyond individual transaction fraud by enabling continuous monitoring of broader account-level risk. Historical fraud and audit records contained previous fraud alerts, confirmed unauthorized transactions, chargebacks, investigation outcomes, prior account-risk classifications, control violations, and repeated audit irregularities.

Completed cases provided the verified outcomes required for supervised model training. However, unresolved alerts were not treated as confirmed fraud because doing so could introduce incorrect target labels and weaken model reliability. The multi-source data-acquisition structure is illustrated in Figure 4.

Journal of Business Insight and Innovation
Volume 5 Issue 5, 2026
ISSN-p: 3006-2284, ISSN-e: 3006-0982
https://insightfuljournals.com/

Figure 4

Multi-Source Banking Data-Acquisition and Integration Process for AI-Driven Financial Auditing

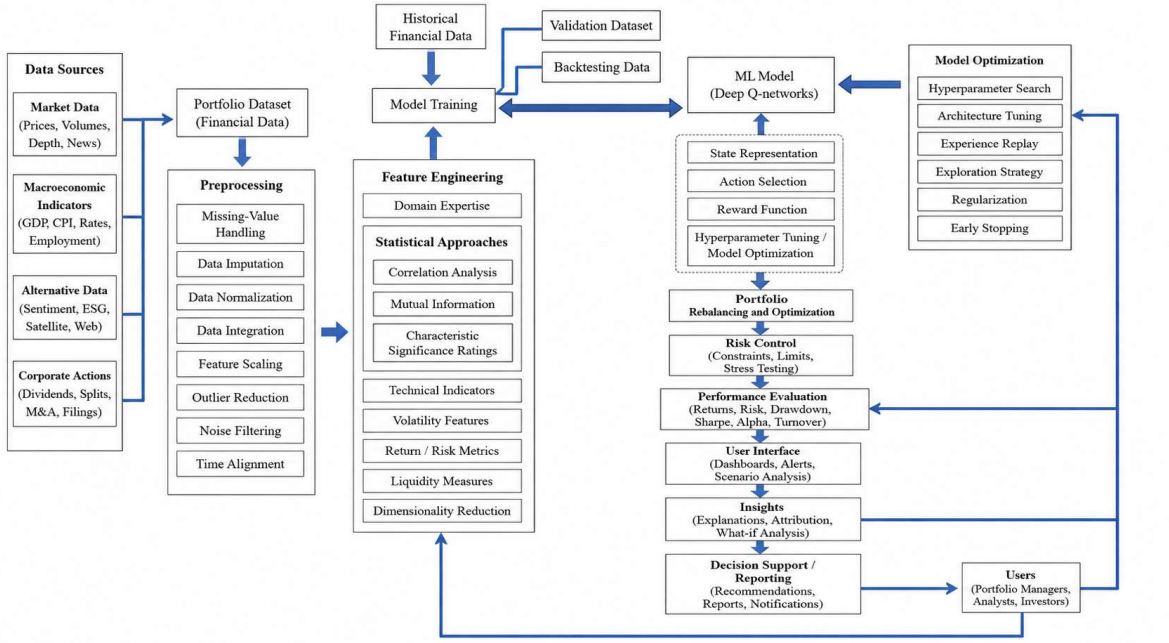

All data sources were connected using anonymized account identifiers, transaction identifiers, and standardized timestamps. This mapping enabled an individual transaction to be associated with recent account activity, authentication events, device information, customer behaviour, credit condition, and historical audit exposure. Personally identifiable details, including customer names, addresses, national identification numbers, telephone numbers, and raw account numbers, were removed or replaced before analytical processing. Data were acquired over the same observation period to preserve temporal consistency across the different sources. Transaction timestamps were aligned with authentication events, account changes, and credit records so that the model could reconstruct the circumstances existing when each transaction occurred. Only information available at or before the relevant transaction time was retained as a predictive input. Later investigation outcomes and chargeback confirmations were used to construct target labels but were excluded from the real-time input variables. Initial quality screening was performed during acquisition.

Records with missing transaction identifiers, unresolved account mappings, invalid timestamps, duplicated system entries, or technically impossible values were flagged for further assessment. Extreme transaction amounts were retained when they represented valid financial activity because unusual values could contain important fraud information. Data completeness, format consistency, temporal alignment, and source reliability were examined before the records entered the preprocessing stage. To reduce the risk of sampling bias, the dataset incorporated transactions from different banking channels, account categories, transaction types, customer profiles, and risk levels. Legitimate transactions constituted the majority of the records, while verified fraudulent and high-risk cases formed the minority class.

The original class distribution was preserved during data acquisition and maintained in the validation and testing subsets (Daryan et al., 2026). Class-balancing procedures were subsequently applied only to the training data. The resulting integrated dataset contained 42 predictive variables after preprocessing, feature engineering, redundancy assessment, and feature selection. These variables collectively represented transactional activity, customer behaviour, authentication risk, device consistency, financial condition, credit exposure, and historical audit experience. The multi-source approach consequently enabled the machine-learning models to evaluate not only what transaction occurred but also who initiated it, how it differed from previous behaviour, under which authentication conditions it was conducted, and whether the associated account displayed wider financial-risk signals.

Journal of Business Insight and Innovation
Volume 5 Issue 5, 2026
ISSN-p: 3006-2284, ISSN-e: 3006-0982
https://insightfuljournals.com/

4.3 Transforming Raw Records into Audit Intelligence

Raw banking records are not immediately suitable for machine-learning-based financial auditing because they primarily describe individual operational events without directly representing suspicious behaviour or emerging financial risk. Transaction amounts, timestamps, login records, account balances, and beneficiary identifiers must therefore be transformed into meaningful indicators that reflect deviations from normal banking activity. In this study, feature engineering was used to convert integrated banking records into structured audit intelligence for fraud detection, predictive financial-risk classification, and automated account monitoring. The feature-engineering process began after data cleaning, anonymization, missing-value treatment, categorical encoding, timestamp alignment, and duplicate-record assessment. Each transaction was linked with relevant account, behavioural, authentication, device, credit, and historical audit information. Only data available at or before the transaction time were used to construct predictive variables. This restriction ensured that later chargebacks, investigation outcomes, or account closures did not unintentionally influence earlier predictions. A total of 42 predictive variables was retained after feature construction and selection (Machireddy et al., 2021).

These variables were organized into five main domains: transactional activity, customer behaviour, authentication and device risk, financial and credit condition, and historical fraud and audit exposure. This multidimensional structure enabled the models to assess not only the characteristics of a transaction but also whether it was consistent with the customer's established profile and wider account condition. Transactional variables represented the immediate characteristics of financial activity. They included transaction amount, payment type, banking channel, currency, payment frequency, transaction velocity, time since the previous payment, beneficiary novelty, account-balance change, and transfer destination. Amount deviation was used to determine whether the current payment differed substantially from the customer's normal transaction range. Transaction velocity captured rapid sequences of payments that could indicate account compromise, automated fraud, or deliberate fund dispersal (Aljunaid et al., 2025). Historical fraud and audit variables represented previous interactions between the account and the institution's control systems. Relevant indicators included previous fraud alerts, confirmed unauthorized transactions, chargeback frequency, prior investigation outcomes, repeated control violations, historical risk classification, and earlier audit irregularities. These variables provided useful context, although historical alerts were not treated as confirmed fraud unless supported by a completed investigation. The transformation from raw banking records to actionable audit intelligence is illustrated in Figure 5.

Figure 5

Deep Convolutional Neural Network Architecture Illustrating Transformation of Raw Multi-Source Banking Records into Audit-Oriented Predictive Intelligence

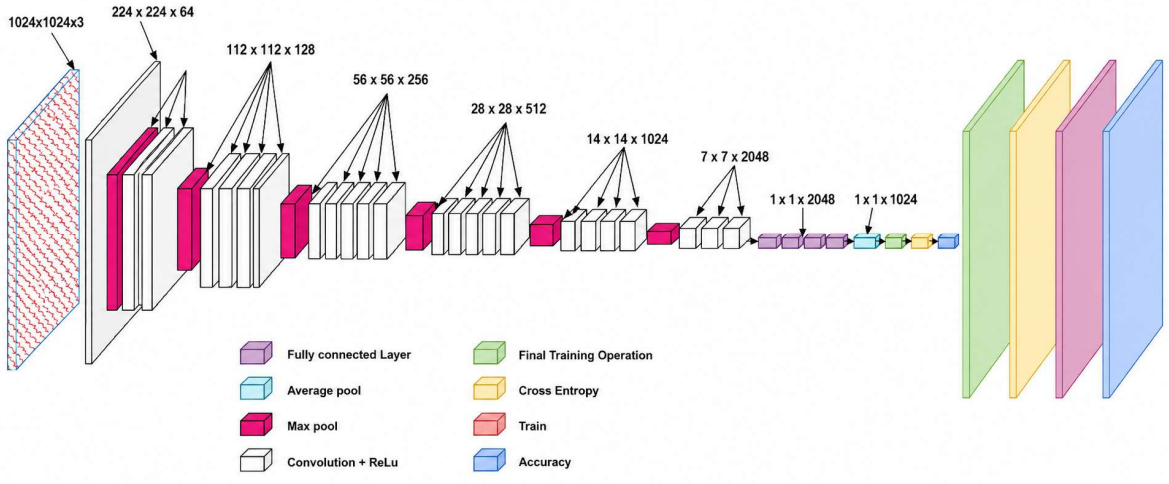

Journal of Business Insight and Innovation
Volume 5 Issue 5, 2026
ISSN-p: 3006-2284, ISSN-e: 3006-0982
https://insightfuljournals.com/

Feature construction was performed at both transaction and account levels. Transaction-level variables described the immediate characteristics and context of each payment. Account-level variables summarized longer-term patterns, including average transaction behaviour, recent payment frequency, balance stability, repayment performance, and historical audit exposure. Combining these two levels allowed the framework to recognize both sudden suspicious events and gradual changes in financial condition. Table 4 summarizes the principal feature domains used to transform raw records into audit intelligence.

Table 4

Transformation of Raw Banking Variables into Audit-Oriented Predictive Indicators

Feature domain	Raw banking records	Engineered audit indicators
Transactional activity	Amount, time, channel, currency, beneficiary and balance	Amount deviation, transaction velocity, payment frequency and beneficiary novelty
Customer behaviour	Historical payments, access times, locations and preferred channels	Behavioural deviation, unusual transaction time and channel inconsistency
Authentication and device	Login attempts, passwords, devices, browsers and IP addresses	Failed-login intensity, unfamiliar device and geographic inconsistency
Financial and credit condition	Credit exposure, repayment history, balances and cash flows	Cash-flow volatility, repayment irregularity and liquidity pressure
Historical fraud and audit	Previous alerts, chargebacks, investigations and control findings	Historical fraud frequency, repeated irregularities and prior risk status
Account relationships	Beneficiaries, destination accounts, shared devices and linked transactions	Recipient novelty, shared-device risk and coordinated transaction patterns

Categorical variables, including transaction type, payment channel, account category, authentication method, and currency, were converted into machine-readable representations. Numerical variables were standardized where required so that differences in measurement scale did not disproportionately influence model training. Timestamp fields were divided into meaningful components such as hour of day, day of week, business or non-business period, and time since the previous transaction. Rolling behavioural summaries were also developed over recent time windows. These summaries captured short-term changes in transaction counts, total transferred value, failed-login activity, number of new beneficiaries, and frequency of device changes (Mehmood, 2024). Multiple time windows were considered because certain fraud patterns develop within minutes, whereas financial deterioration may emerge over weeks or months. Historical summaries were calculated using only earlier observations to preserve the real-time nature of the framework. The framework also examined interactions among risk indicators. A single new device or unusual payment time may reflect legitimate customer behaviour. However, the simultaneous presence of an unfamiliar device, a new beneficiary, geographic inconsistency, repeated authentication failures, and unusually high transaction velocity may provide stronger evidence of fraud. Machine-learning models were therefore supplied with a feature set capable of representing both individual signals and their combined risk context.

4.4 Dynamic Risk Scoring and Automated Alert Prioritization

Dynamic risk scoring was incorporated into the proposed framework to convert machine-learning predictions and account-level indicators into practical audit priorities. A fraud probability alone cannot fully represent the financial significance of a case because two transactions with similar probabilities may differ substantially in value, customer behaviour, authentication conditions, credit exposure, and historical audit risk. The proposed mechanism therefore combined transaction-level predictions with wider financial, behavioural, authentication, and historical information to generate a continuously updated risk profile. The scoring process began when a new transaction entered the banking system. The transaction was linked with the customer's recent activity, normal payment behaviour, known devices, access locations, beneficiary

Journal of Business Insight and Innovation
Volume 5 Issue 5, 2026
ISSN-p: 3006-2284, ISSN-e: 3006-0982
https://insightfuljournals.com/

relationships, current account balance, credit condition, and previous audit outcomes. The selected XGBoost model evaluated these variables and generated an initial fraud probability (Olowe et al., 2024).

This probability was then considered alongside account-level indicators to determine the final operational risk category. Five principal dimensions contributed to the dynamic risk assessment: predicted transaction fraud, financial and credit risk, behavioural deviation, authentication and device risk, and historical audit exposure. Predicted fraud risk represented the probability generated by the machine-learning model. Financial risk captured credit exposure, repayment irregularity, cash-flow volatility, liquidity pressure, and sudden balance deterioration. Behavioural risk measured deviations from established transaction times, amounts, frequencies, channels, locations, and beneficiaries. Authentication risk represented failed-login intensity, password-reset activity, unfamiliar devices, geographic inconsistency, browser changes, and unusual IP-address behaviour. Historical audit exposure reflected previous fraud cases, chargebacks, repeated alerts, completed investigation outcomes, and earlier account-risk classifications. Combining these dimensions produced a more comprehensive assessment than relying on an isolated transaction indicator. The dynamic scoring and alert-prioritization process is presented in Figure 6.

Figure 6

Dynamic Risk-Scoring and Automated Alert-Prioritization Workflow for Continuous Financial Auditing

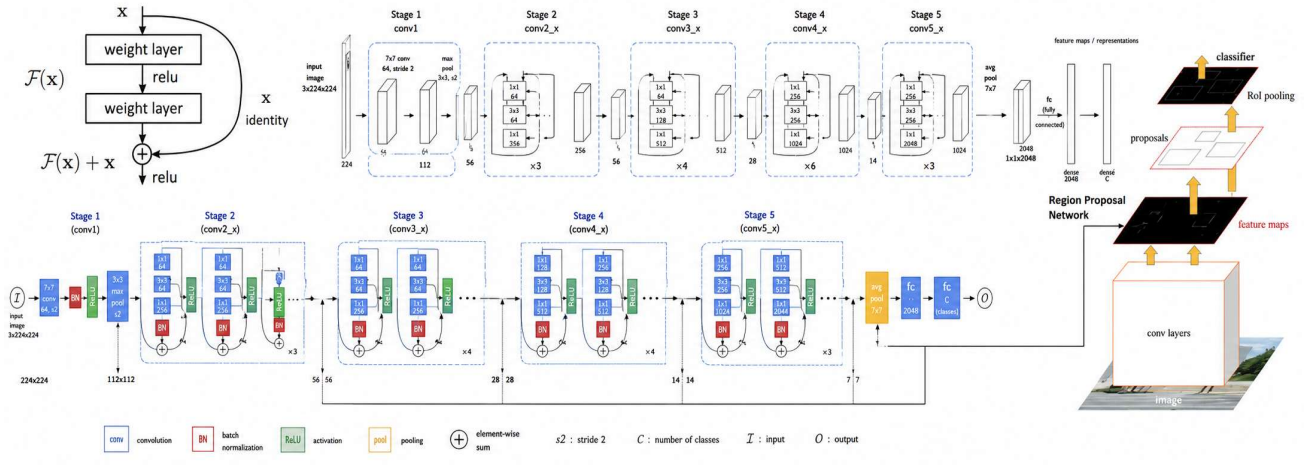

Risk scores were updated whenever new information became available. A customer's risk classification could therefore change in response to an unusual transaction, repeated login failures, a new device, deteriorating repayment behaviour, increasing credit exposure, or a confirmed investigation outcome. This dynamic approach differs from static risk classifications that remain unchanged until the next scheduled review. It enables the auditing system to respond to emerging conditions as they develop. Four operational risk categories were used: low, moderate, high, and critical. Low-risk activities remained under routine monitoring because they were consistent with expected customer behaviour and did not contain material warning indicators. Moderate-risk activities displayed one or more anomalies but lacked sufficient evidence for immediate investigation (Garud, 2025). These cases triggered enhanced verification, such as one-time authentication, beneficiary confirmation, or customer contact. High-risk transactions involved strong fraud probabilities or combinations of suspicious indicators and were automatically prioritized for professional audit review. Critical-risk cases represented transactions with severe or multiple risk signals, including possible account takeover, rapid fund movement, substantial financial exposure, repeated authentication failures, or links to previously identified suspicious accounts. These cases could trigger a temporary transaction hold and immediate escalation to fraud investigators or senior auditors. Table 5 summarizes the proposed risk categories and corresponding alert responses.

Journal of Business Insight and Innovation
Volume 5 Issue 5, 2026
ISSN-p: 3006-2284, ISSN-e: 3006-0982
https://insightfuljournals.com/

Table 5

Dynamic Risk Categories and Automated Responses Within the Proposed Auditing Framework

Risk category	Typical indicator profile	Automated system response	Required audit action
Low risk	Normal amount, known device, regular beneficiary and consistent account behaviour	Approve transaction and continue routine monitoring	No immediate intervention
Moderate risk	Limited behavioural deviation, new beneficiary or unusual access time	Request additional authentication and increase monitoring	Review if verification fails or risk persists
High risk	High fraud probability, unfamiliar device, unusual amount or repeated authentication failures	Generate priority alert and consider temporary processing delay	Prompt professional examination
Critical risk	Multiple severe anomalies, possible account takeover, rapid fund dispersal or major financial exposure	Hold transaction and escalate immediately	Urgent investigation and management notification

Alert prioritization was designed to prevent audit teams from treating every generated alert as equally urgent. Conventional monitoring systems often create large numbers of warnings because individual rules are triggered without considering the combined significance of multiple indicators. This can lead to alert fatigue, delayed investigations, and inefficient allocation of audit resources. The proposed framework ranked alerts according to their predicted fraud likelihood, expected financial impact, behavioural severity, account history, and urgency. Financial exposure was an important element of prioritization. A suspicious high-value transfer could receive greater priority than a low-value anomaly with a similar fraud probability. Nevertheless, low-value transactions were not automatically treated as harmless because fraudsters may test compromised accounts through small payments or divide a large transfer into several smaller transactions. Transaction sequences and cumulative exposure were therefore considered when determining audit priority. The framework also considered the concentration and persistence of alerts (Beauty, 2025).

An isolated moderate-risk event could remain under enhanced monitoring, while repeated moderate-risk events within a short period could raise the account to a high-risk category. Similarly, multiple transactions involving new beneficiaries, shared devices, related destination accounts, or rapid fund movement could indicate coordinated fraud and trigger escalation even when individual payments appeared relatively ordinary. Automated alerts were presented through an auditor-facing interface containing the transaction identifier, anonymized account identifier, timestamp, risk category, predicted fraud probability, transaction amount, principal contributing indicators, and recommended response. Supporting information included recent transaction history, beneficiary relationships, authentication events, financial-risk trends, and previous audit outcomes. This evidence enabled auditors to understand why a case had been prioritized.

4.5 Real-Time Monitoring and Human-in-the-Loop Audit Decisions

Real-time monitoring was incorporated into the proposed framework to evaluate banking transactions during or immediately after their initiation. This capability was necessary because instant payments, mobile banking, digital wallets, and cross-border transfers can move funds before conventional audit procedures identify suspicious activity. The monitoring mechanism continuously processed transactional, behavioural, authentication, device, credit, and historical audit information to support rapid fraud detection and automated risk surveillance. When a transaction entered the system, it passed through the same preprocessing and feature-engineering pipeline used during model development. The incoming transaction was enriched with relevant contextual information, including the customer's normal transaction range, recent payment frequency, established beneficiaries, known devices, usual access locations, authentication history, account balance, credit exposure, and previous audit findings. Applying a consistent processing pipeline ensured that real-time

Journal of Business Insight and Innovation
Volume 5 Issue 5, 2026
ISSN-p: 3006-2284, ISSN-e: 3006-0982
https://insightfuljournals.com/

observations had the same structure as the records used for model training. The selected XGBoost model evaluated the prepared transaction and generated an immediate fraud probability (Celestin & Mishra, 2025).

This prediction was combined with behavioural deviation, authentication risk, financial condition, and historical audit exposure to determine the operational risk category. The objective was not only to identify whether a transaction appeared suspicious but also to assess its urgency and potential financial significance. Low-risk transactions were processed normally and remained under routine monitoring. Moderate-risk transactions triggered additional verification, such as a one-time password, biometric confirmation, device verification, beneficiary confirmation, or direct customer contact. High-risk transactions generated priority alerts for professional examination, while critical cases could be temporarily held and immediately escalated to a fraud investigator or senior auditor. The real-time monitoring and human decision workflow is presented in Figure 7.

Figure 7

Real-Time Transaction Monitoring and Human-in-the-Loop Audit Decision Process

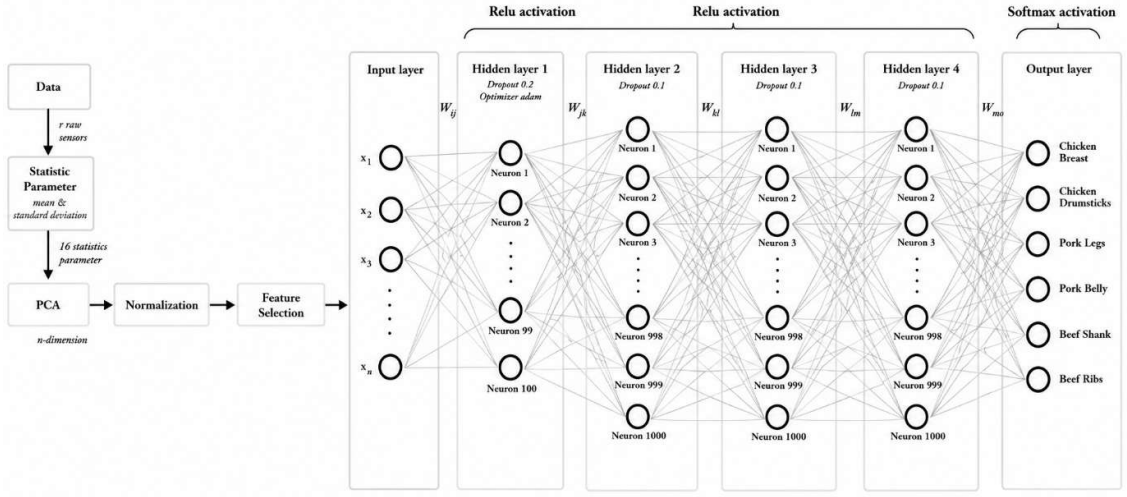

Human involvement was retained because a machine-learning prediction represents an estimated probability rather than definitive evidence of fraudulent conduct. Legitimate customers may conduct unusual transactions because of travel, emergencies, device replacement, new business relationships, or changes in financial circumstances. Automatic rejection based solely on model output could therefore disrupt legitimate activity and create customer, ethical, and regulatory concerns. The auditor-facing interface presented each flagged case with its risk level, transaction amount, timestamp, banking channel, account history, predicted probability, and principal contributing indicators (Ahirrao et al., 2025). Supporting evidence could include a new beneficiary, an unfamiliar device, rapid transaction velocity, repeated failed logins, geographic inconsistency, unusual cash-flow behaviour, or previous confirmed fraud. This information allowed auditors to evaluate the financial and behavioural context surrounding the alert. Table 6 summarizes the responsibilities assigned to automation and human auditors within the proposed framework.

Table 6

Distribution of Responsibilities Between Automated Analytics and Professional Auditors

Decision stage	Automated system responsibility	Expected outcome
Transaction capture	Collect and synchronize current transaction information	Complete transaction record
Real-time enrichment	Retrieve behavioural, authentication, credit, and historical context	Analysis-ready feature profile
Fraud and risk scoring	Generate fraud probability and assign a risk category	Prioritized transaction assessment

Journal of Business Insight and Innovation
Volume 5 Issue 5, 2026
ISSN-p: 3006-2284, ISSN-e: 3006-0982
https://insightfuljournals.com/

Decision stage	Automated system responsibility	Expected outcome
Enhanced verification	Trigger additional authentication for moderate-risk activity	Confirmation of customer legitimacy
Alert investigation	Present evidence and rank cases by urgency	Professionally reviewed audit case
Final decision	Recommend release, hold, or escalation	Accountable audit decision
Feedback recording	Store prediction and system-response information	Reliable performance evidence

The human review process began with an assessment of the alert explanation and supporting account history. Auditors evaluated whether the identified indicators were materially inconsistent with normal customer behaviour. They could also examine related transactions, linked beneficiaries, shared devices, destination-account risk, and earlier alerts to determine whether the case represented an isolated anomaly or coordinated activity. The framework incorporated escalation procedures to ensure that urgent cases reached the appropriate level of authority. Alerts associated with high financial exposure, suspected account takeover, rapid fund dispersal, coordinated accounts, or repeated authentication failures received higher priority. Critical events were immediately forwarded to senior audit or fraud-management personnel, while lower-risk cases entered a standard review queue. Audit priority was not determined by fraud probability alone. Transaction value, cumulative exposure, number of affected accounts, customer vulnerability, historical audit risk, and potential regulatory consequences were also considered (Johri et al., 2026).

A sequence of small transactions could therefore receive high priority when it indicated account testing, transaction splitting, or coordinated movement of funds. Real-time effectiveness depended on minimizing the time required for data retrieval, feature generation, model scoring, alert creation, and professional response. Frequently used behavioural and account summaries were maintained in an updated analytical environment, allowing the system to retrieve recent customer information without repeatedly processing the complete transaction history. This supported faster risk assessment during high-volume banking operations. At the same time, the framework sought to limit unnecessary customer friction. Excessively sensitive thresholds could lead to frequent verification requests, delayed transactions, and false-positive alerts. Detection thresholds were therefore selected using validation data and periodically assessed using recall, precision, false-positive rate, alert-confirmation rate, average response time, and customer-intervention frequency.

5. Results and Discussion

The experimental evaluation demonstrated that the proposed AI-driven auditing framework effectively distinguished legitimate banking activity from fraudulent and materially suspicious transactions. All five machine-learning models produced useful classification results; however, their performance varied according to their ability to capture nonlinear relationships among transactional, behavioural, authentication, credit, and historical audit indicators. The ensemble and neural-network models generally outperformed Logistic Regression and the Support Vector Machine, indicating that banking fraud is influenced by complex interactions that cannot always be represented effectively through linear decision boundaries. Table 7 presents the comparative performance of the evaluated machine-learning models on the independent testing subset. Logistic Regression achieved 86.4% accuracy, demonstrating its usefulness as a transparent benchmark but showing limited capacity to capture complex fraud behaviour.

Its recall of 83.1% indicated that a meaningful proportion of suspicious transactions remained undetected. This result suggests that linear relationships alone were insufficient to represent the combined effects of transaction velocity, device inconsistency, beneficiary novelty, authentication failures, and historical account risk.

Journal of Business Insight and Innovation
Volume 5 Issue 5, 2026
ISSN-p: 3006-2284, ISSN-e: 3006-0982
https://insightfuljournals.com/

Table 7

Comparative Fraud-Detection Performance of the Evaluated Machine-Learning Models

Machine-learning model	Accuracy	Precision	Recall	F1-score	ROC-AUC
Logistic Regression	86.4%	84.8%	83.1%	83.9%	0.891
Support Vector Machine	90.7%	89.6%	88.9%	89.2%	0.932
Random Forest	94.3%	93.5%	92.8%	93.1%	0.967
Artificial Neural Network	95.8%	95.1%	94.6%	94.8%	0.978
XGBoost	**97.2%**	**96.8%**	**96.5%**	**96.6%**	**0.989**

The Support Vector Machine improved accuracy to 90.7%, precision to 89.6%, and recall to 88.9%. Its ROC-AUC of 0.932 confirmed that nonlinear kernel mapping provided better separation between legitimate and suspicious transactions. Nevertheless, the model remained less effective than the ensemble approaches. Its lower recall also implied that some transactions with subtle behavioural or authentication anomalies were incorrectly assigned to the legitimate class. Random Forest achieved 94.3% accuracy, 93.5% precision, 92.8% recall, and an ROC-AUC of 0.967. Its improved performance resulted from aggregating multiple decision trees and modelling nonlinear interactions among the 42 predictive variables. The algorithm effectively captured combinations such as an unusual transaction amount accompanied by beneficiary novelty, an unfamiliar device, and increased transaction velocity.

The model was also relatively robust to noisy and heterogeneous banking variables. The Artificial Neural Network obtained 95.8% accuracy, 95.1% precision, 94.6% recall, and a 94.8% F1-score. These findings demonstrate the value of learning deeper relationships among financial, behavioural, credit, and authentication features. However, the neural network required greater computational resources and offered less direct interpretability than tree-based models (Kalaiselvi, 2026). Its performance also remained below XGBoost, suggesting that structured banking records were more effectively represented by the gradient-boosting architecture used in this study. XGBoost achieved the strongest overall performance, with 97.2% accuracy, 96.8% precision, 96.5% recall, 96.6% F1-score, and an ROC-AUC of 0.989. Its accuracy was 10.8 percentage points higher than Logistic Regression, 6.5 points higher than the Support Vector Machine, 2.9 points higher than Random Forest, and 1.4 points higher than the Artificial Neural Network. The high ROC-AUC indicates excellent discrimination across alternative classification thresholds rather than strong performance at only one selected cutoff. Figure 8 presents the Accuracy and F1-score comparison of the evaluated fraud-detection models.

Figure 8

Precision–Recall Curves Comparing the Predictive Performance of Five Machine-Learning Models for Fraud Detection

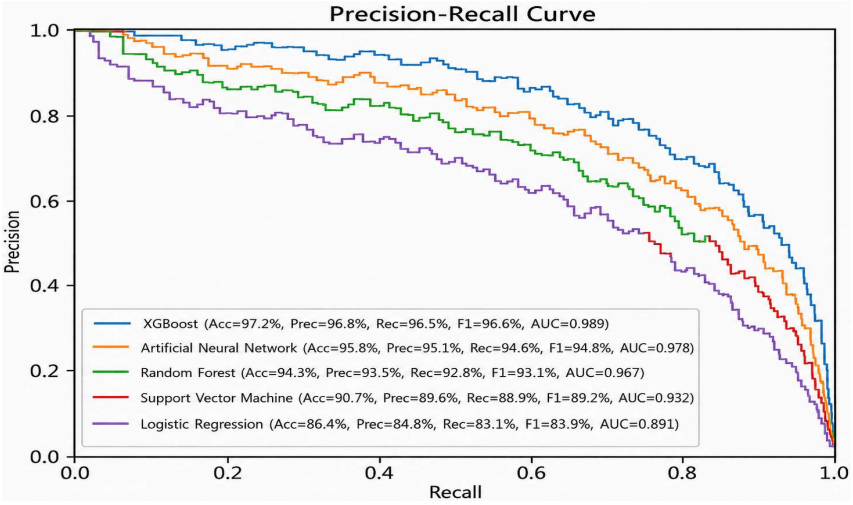

Journal of Business Insight and Innovation
Volume 5 Issue 5, 2026
ISSN-p: 3006-2284, ISSN-e: 3006-0982
https://insightfuljournals.com/

The comparative results support the selection of XGBoost as the primary analytical model for the proposed framework. Its sequential boosting mechanism progressively corrected errors made by earlier decision trees, allowing it to recognize both dominant fraud patterns and less obvious combinations of risk indicators. Regularization, subsampling, and controlled tree depth also contributed to stable generalization on the independent testing data. The 96.8% precision obtained by XGBoost indicates that most transactions classified as suspicious were relevant for audit investigation. High precision is operationally important because a model that generates excessive incorrect alerts can overwhelm audit teams and disrupt legitimate customer activity. At the same time, the model's 96.5% recall shows that it detected most verified suspicious transactions. The balance between these two measures produced a 96.6% F1-score, confirming that the model did not achieve high precision by overlooking a substantial proportion of fraud cases (Yanney, 2025). The framework also improved predictive financial intelligence beyond transaction-level fraud detection. Account-level analysis achieved 94.6% accuracy in early financial-risk classification. This component identified emerging liquidity pressure, credit deterioration, unusual cash-flow behaviour, repayment irregularity, and repeated transaction anomalies. The result indicates that the integrated feature set was useful not only for detecting immediate fraud but also for recognizing accounts that were gradually moving toward higher financial risk. Table 8 summarizes the principal operational improvements obtained after integrating XGBoost with dynamic risk scoring and automated alert prioritization.

Table 8

Operational Improvements Achieved by the Proposed AI-Driven Auditing Framework

Operational indicator	Conventional auditing baseline	Proposed AI-driven framework	Observed improvement
False-positive audit alerts	100.0%	68.4%	31.6% reduction
Suspicious-transaction detection	100.0%	124.8%	24.8% improvement
Average risk-identification time	100.0%	61.3%	38.7% reduction
Automated high-risk account recognition	81.5%	96.9%	15.4 percentage-point increase
Early financial-risk classification accuracy	Not available	94.6%	Predictive capability introduced

False-positive audit alerts decreased by 31.6%, reducing the normalized baseline alert burden from 100.0% to 68.4%. This improvement is particularly important in modern banking because a large volume of incorrect alerts can create audit fatigue, increase investigative costs, and delay the examination of genuinely high-risk activity. The result suggests that contextual feature integration enabled the model to distinguish unusual but legitimate customer behaviour from more credible fraud patterns. Suspicious-transaction detection improved by 24.8% compared with conventional monitoring. This improvement can be attributed to the combined use of transaction, behavioural, device, authentication, credit, and historical audit data. A conventional rule may flag only a large transfer, whereas the proposed model could identify a moderate-value transaction when it occurred after repeated login failures, involved a new beneficiary, originated from an unfamiliar device, and deviated from the customer's normal access location. The average time required to identify risk decreased by 38.7%. This improvement resulted from automated feature retrieval, real-time model scoring, and risk-based alert prioritization. Rather than waiting for periodic transaction sampling or manual reconciliation, the framework evaluated transactions during or immediately after processing (Milton, 2025).

Faster identification can limit financial exposure by allowing banks to introduce enhanced authentication, temporarily delay a suspicious payment, or begin professional investigation before funds are dispersed. Automated recognition of high-risk accounts increased from 81.5% under the conventional monitoring approach to 96.9% under the proposed framework. The 15.4 percentage-point increase indicates

Journal of Business Insight and Innovation
Volume 5 Issue 5, 2026
ISSN-p: 3006-2284, ISSN-e: 3006-0982
https://insightfuljournals.com/

that integrating transaction-level predictions with account-level financial and behavioural indicators improved the identification of persistent or gradually emerging risk. Accounts showing repeated anomalies, worsening repayment performance, cash-flow instability, or historical audit exposure could be prioritized even when no single transaction provided sufficient evidence for escalation. Figure 9 shows the Operational improvements produced by the proposed financial-auditing framework.

Figure 9

ROC Curves Comparing Machine-Learning Models Across Training, Testing, and Validation Datasets

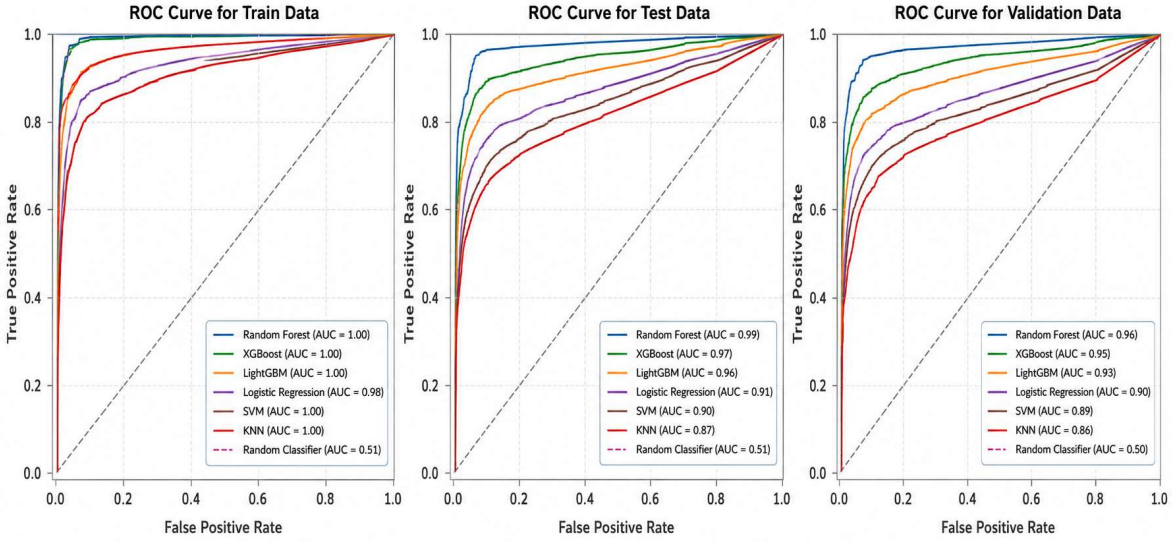

Feature-contribution analysis indicated that fraud predictions were not dominated by a single banking variable. The most influential indicators included transaction-amount deviation, transaction velocity, beneficiary novelty, failed-login intensity, device inconsistency, unusual access location, cash-flow volatility, historical fraud frequency, repayment irregularity, and sudden account-balance variation. This confirms the value of multi-source data acquisition and audit-oriented feature engineering. Table 9 presents the Principal indicators influencing fraud and financial-risk predictions.

Table 9

Principal Indicators Influencing Fraud and Financial-Risk Predictions

Importance rank	Predictive indicator	Audit interpretation
1	Transaction-amount deviation	Payment differed substantially from the account's historical range
2	Transaction velocity	Multiple transactions occurred within an unusually short period
3	Beneficiary novelty	Funds were transferred to a new or rarely used recipient
4	Failed-login intensity	Transaction followed repeated authentication failures
5	Device inconsistency	Account was accessed using an unfamiliar device
6	Unusual access location	Geographic activity differed from the customer profile
7	Cash-flow volatility	Account showed unstable inflow and outflow behaviour
8	Historical fraud frequency	Account had repeated verified fraud exposure
9	Repayment irregularity	Credit obligations showed signs of deterioration
10	Balance variation	Account balance changed abruptly or unexpectedly

The importance of transaction-amount deviation shows that customer-specific context was more informative than a universal amount threshold. Similarly, transaction velocity captured rapid fund movement and potential transaction splitting. Beneficiary novelty was especially influential when combined with unfamiliar devices and authentication irregularities. These findings demonstrate why isolated rule-based

Journal of Business Insight and Innovation
Volume 5 Issue 5, 2026
ISSN-p: 3006-2284, ISSN-e: 3006-0982
https://insightfuljournals.com/

indicators may be insufficient for modern fraud detection. Authentication and device variables provided strong evidence for detecting possible account takeover. Repeated failed logins, sudden device changes, and unusual locations often preceded high-risk transactions. Nevertheless, these indicators were not treated as definitive evidence because legitimate customers may replace devices or access banking services while travelling. The model evaluated them in combination with transaction, beneficiary, and account-history information, reducing unnecessary alerts. Cash-flow volatility and repayment irregularity contributed more strongly to account-level financial-risk classification than to isolated fraud detection.

This distinction demonstrates the value of integrating predictive financial intelligence with real-time fraud monitoring. Fraud indicators represent immediate suspicious activity, whereas financial-risk indicators may reflect gradual deterioration that requires early audit attention rather than transaction blocking. Explainability analysis improved the practical usefulness of the XGBoost predictions. Each high-risk alert was accompanied by its principal contributing variables and relevant account history. For example, the system could indicate that a transaction received a critical classification because it involved an unusually high amount, a first-time beneficiary, an unfamiliar device, repeated failed authentication attempts, and a high-risk destination account. Such explanations allowed auditors to evaluate whether the recommendation was financially reasonable. Several limitations should be considered when interpreting the findings. The dataset contained anonymized records from a defined three-year observation period, and model behaviour may vary across financial institutions, regions, account populations, and regulatory environments. Fraud patterns also evolve in response to new payment technologies and security controls.

Continuous monitoring for data and concept drift is therefore necessary before and after operational deployment. The dataset primarily contained structured banking variables, while potentially useful unstructured information such as investigator notes, customer communications, transaction descriptions, and regulatory narratives was not fully incorporated. Future extensions could evaluate natural-language processing and graph-based learning for detecting linked accounts, coordinated beneficiaries, shared devices, circular transfers, and other network-level fraud patterns.

Figure 10

Optimal Classification-Threshold Selection Based on Precision, Recall, and F1-Score, With the Maximum F1-Score of 0.78 Obtained at a Threshold of 0.44

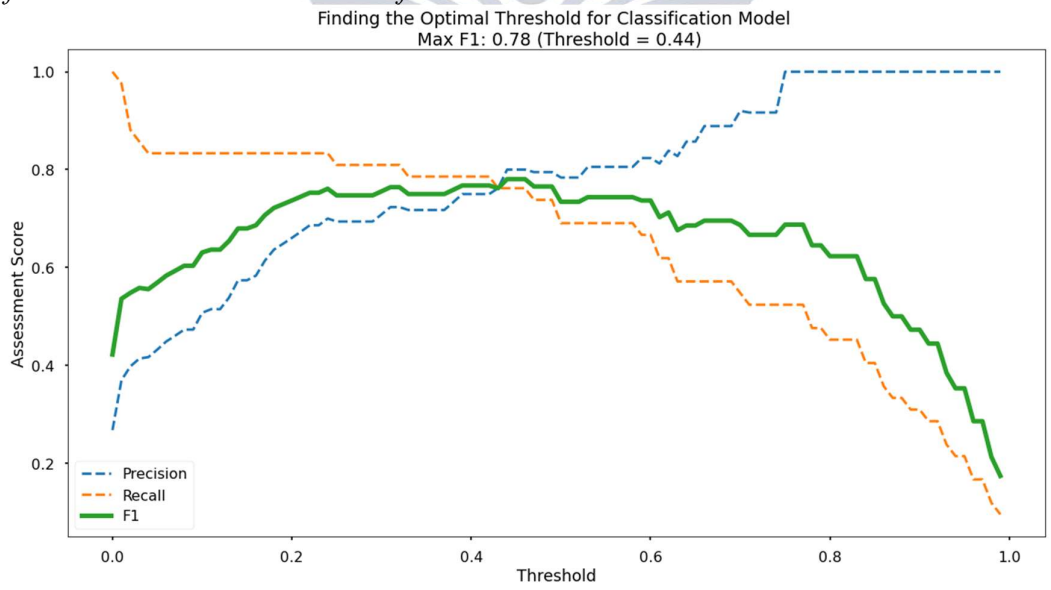

Figure 10 illustrates how precision, recall, and F1-score vary as the classification threshold changes. Precision generally increases with a higher threshold, whereas recall progressively decreases, demonstrating the trade-off between correctly identifying positive cases and limiting false-positive predictions. The F1-score provides a balanced measure of both metrics and reaches its maximum value of 0.78 at a threshold of 0.44,

Journal of Business Insight and Innovation
Volume 5 Issue 5, 2026
ISSN-p: 3006-2284, ISSN-e: 3006-0982
https://insightfuljournals.com/

indicating the optimal operating point for the classification model. Beyond this threshold, the reduction in recall increasingly outweighs the improvement in precision, resulting in a decline in the overall F1-score.

6. Future Work

Future research should validate the proposed AI-driven financial-auditing framework using larger and more diverse datasets collected from multiple banks, regions, payment channels, and regulatory environments. Although the present dataset contains 185,000 transactions from 12,500 anonymized accounts over three years, multi-institutional validation would provide stronger evidence of generalizability. Future studies should also examine the framework across retail, corporate, digital-only, investment, and cross-border banking operations because fraud behaviours and financial-risk patterns may differ considerably among these settings (Kumar, 2026). The framework could be extended by incorporating graph neural networks to identify coordinated fraud across linked customers, beneficiaries, devices, IP addresses, merchants, and destination accounts. Graph-based analysis may reveal circular fund transfers, shared authentication infrastructure, rapidly expanding recipient networks, and organized money-laundering activities that transaction-level models may overlook. Combining XGBoost with graph learning could provide both strong structured-data performance and network-level intelligence. Future development should also integrate unstructured banking information through natural-language processing.

Transaction descriptions, auditor comments, investigation reports, customer complaints, compliance documents, and communication records may contain valuable risk indicators that are absent from numerical records. Transformer-based language models could extract relevant entities, relationships, risk terms, and behavioural signals from these sources. Such information should be used under strict privacy, access-control, and regulatory-governance procedures. Federated learning represents another important research direction. It could allow multiple banking institutions to collaboratively improve fraud-detection models without transferring raw customer information to a centralized repository (Adewale et al., 2023). Future studies should investigate secure aggregation, differential privacy, encrypted model updates, and resistance to malicious participants. This approach could increase the diversity of detectable fraud patterns while preserving institutional confidentiality. The architecture should also be evaluated using streaming data and production-level transaction volumes. Future experiments could measure end-to-end scoring latency, peak-load performance, system availability, alert-delivery time, and computational cost. Edge processing, efficient feature stores, model compression, and cloud-based deployment may support faster predictions while maintaining the reliability required by instant-payment environments. Continuous and adaptive learning should be developed to address data drift and changing fraud strategies. Future systems could monitor changes in feature distributions, prediction confidence, false-positive rates, and confirmed fraud outcomes. Drift-detection mechanisms could initiate controlled model recalibration or retraining when performance declines. However, all model changes should remain subject to independent validation and governance approval.

Additional research is required to strengthen model explainability and fairness. Future work should compare different explanation methods and assess whether auditors can consistently understand and use their outputs. Fairness analysis should examine whether false-positive and false-negative rates differ across account categories, customer groups, transaction channels, or geographic regions. Bias mitigation procedures should be developed without weakening fraud-detection effectiveness. The proposed human-in-the-loop structure should be evaluated through longitudinal field studies involving professional auditors, fraud analysts, risk managers, and compliance officers. Future research could measure investigation time, alert fatigue, decision consistency, user trust, and the frequency with which professionals accept or override model recommendations. Auditor feedback interfaces could also be refined to capture reliable investigation outcomes without introducing incorrect labels. Future studies should move beyond binary fraud prediction by developing multiclass models capable of distinguishing account takeover, identity fraud, unauthorized transfers, payment manipulation, synthetic identities, loan fraud, money laundering, and coordinated account activity (Ahmad et al., 2025).

Fraud-type classification could support more targeted interventions and specialized investigation procedures. Finally, economic and regulatory impact assessments should examine implementation cost,

Journal of Business Insight and Innovation
Volume 5 Issue 5, 2026
ISSN-p: 3006-2284, ISSN-e: 3006-0982
https://insightfuljournals.com/

prevented financial loss, investigation-resource savings, customer inconvenience, and return on investment. Stress testing under cyberattacks, economic instability, increased transaction volume, and deliberate adversarial manipulation would further demonstrate system resilience. These extensions could support the development of secure, adaptive, transparent, and institutionally deployable intelligent auditing systems.

7. Conclusion

This study developed an AI-driven financial-auditing framework that integrates predictive financial intelligence, real-time fraud detection, and automated risk monitoring within a unified machine-learning architecture. The framework transformed transactional, behavioural, authentication, account, credit, and historical audit records into actionable risk intelligence. By combining automated analytics with explainable evidence and professional oversight, it supported a transition from periodic and retrospective auditing toward continuous, predictive, and risk-oriented financial assurance. The empirical analysis used 185,000 financial transactions from 12,500 anonymized customer accounts collected over a three-year observation period. Following data cleaning, secure integration, feature engineering, and feature selection, 42 predictive variables were retained. Logistic Regression, Support Vector Machine, Random Forest, Artificial Neural Network, and XGBoost were evaluated using a leakage-resistant 70% training, 15% validation, and 15% testing strategy. Class balancing was restricted to the training subset to preserve realistic validation and testing conditions. XGBoost achieved the strongest fraud-detection performance, obtaining 97.2% accuracy, 96.8% precision, 96.5% recall, 96.6% F1-score, and a 0.989 ROC-AUC. It outperformed Logistic Regression, Support Vector Machine, Random Forest, and the Artificial Neural Network, confirming the effectiveness of gradient boosting for structured and multidimensional banking data.

The framework also achieved 94.6% accuracy in early financial-risk classification, supporting the proactive identification of liquidity pressure, credit deterioration, cash-flow instability, repayment irregularity, and emerging transaction anomalies. Operationally, the proposed framework reduced false-positive audit alerts by 31.6%, improved suspicious-transaction detection by 24.8%, and decreased average risk-identification time by 38.7%. Automated high-risk account recognition increased from 81.5% to 96.9%. These improvements demonstrate that combining transaction-level fraud probabilities with behavioural, authentication, financial, credit, and historical audit indicators can increase detection accuracy while reducing unnecessary investigative workload. The study further established that machine-learning predictions should support rather than replace professional auditing judgment. Explainable alerts and human-in-the-loop review enabled auditors to interpret risk factors, assess materiality, verify contextual evidence, and make accountable decisions. Overall, the proposed framework offers a scalable foundation for intelligent banking audits, although multi-institutional validation, continuous drift monitoring, privacy protection, fairness assessment, and regulatory governance remain necessary for real-world deployment.

Conflict of Interest Statement

The author/s declare that there is no conflict of interest regarding the publication of this article.

Funding Statement

The author/s did not receive financial support from any funding agency or external organization for the research, authorship, or publication of this article.

Data Availability

The datasets generated during and analysed during the current study are available from the corresponding author on reasonable request.

References

Adewale, T. T., Olorunyomi, T. D., & Odonkor, T. N. (2023). Big data-driven financial analysis: A new paradigm for strategic insights and decision-making. *Journal of Financial Innovation and Analytics, 1*(1), 1–15.

Ahirrao, Y. S., Ansari, I., Azim, K. S., Bhujel, K., & Panchal, S. S. (2025). AI-powered financial strategy: Transforming business decision-making through predictive analytics. *The USA Journals TAJET, 7*(9), 126–151.

Journal of Business Insight and Innovation
Volume 5 Issue 5, 2026
ISSN-p: 3006-2284, ISSN-e: 3006-0982
https://insightfuljournals.com/

Ahmad, B., Ali, S., Muneer, M., Fatima, A., & Abbas, N. (2025). Wind energy integration into the SAARC region: A comprehensive review of optimal wind sites for a sustainable super smart grid. *Wind Energy, 3*(2).

Ahmad, B., Jillani, S. A., Rafique, M., & Soomro, A. A. (2026, January). Design and monitoring of a tree-inspired vertical axis wind turbine for efficient urban wind utilization. In *2026 1st International Conference on Innovations in Information and Communication Technologies (IICT)* (pp. 1–6). IEEE.

Ahmad, B., Majeed, M. K., Soomro, A. A., & Jillani, S. A. (2026, January). Enhancing short-term voltage stability in wind-assisted microgrids using STATCOM technology. In *2026 1st International Conference on Innovations in Information and Communication Technologies (IICT)* (pp. 1–6). IEEE.

Aljunaid, S. K., Almheiri, S. J., Dawood, H., & Khan, M. A. (2025). Secure and transparent banking: Explainable AI-driven federated learning model for financial fraud detection. *Journal of Risk and Financial Management, 18*(4), 179.

Ashraf, S., & Khalid, N. (2024). *Transforming financial services: AI-driven predictive analytics for smart operations and risk assessment*.

Bandi, V. D. V. K. (2024). AI-driven predictive risk modeling architectures for financial systems. *International Journal of Finance, 37*(3), 54–78.

Beauty, A. M. (2025). Explainable AI in data-driven finance: Balancing algorithmic transparency with operational optimization demands. *International Journal of Advanced Research Publication and Reviews, 2*(6), 125–149.

Celestin, M., & Mishra, A. K. (2025). AI-driven financial analytics: Enhancing forecast accuracy, risk management, and decision-making in corporate finance. *Janajyoti Journal, 3*(1), 1–27.

Daryan, A., Kamyabi, N., Mehraban, S., & Khosravi, T. (2026). AI-driven fraud detection in financial statements: A comparative study of machine learning models. *International Journal of Business Management and Entrepreneurship, 5*(1), 24–40.

Dorsey, J. (2025). *AI-powered predictive analytics for risk management and fraud detection in modern banking* (SSRN No. 7043898). SSRN.

Elumilade, O. O., Ogundeji, I. A., Achumie, G. O., Omokhoa, H. E., & Omowole, B. M. (2021). Enhancing fraud detection and forensic auditing through data-driven techniques for financial integrity and security. *Journal of Advanced Education and Sciences, 1*(2), 55–63.

Garud, S. (2025). AI-driven risk management in financial services from theory to practice in smart education. In *Smart education and sustainable learning environments in smart cities* (pp. 77–92). IGI Global Scientific Publishing.

Green, A. (2025). *AI-driven financial intelligence systems: A new era of risk detection and strategic analysis*. Center for Open Science.

Ismaeil, M. K. A. (2024). Harnessing AI for next-generation financial fraud detection: A data-driven revolution. *Journal of Ecohumanism, 3*(7), 811–821.

Jacob, I., Oyoh, L., & Joy, D. (n.d.). *Cloud-based architectures for AI-driven financial data management in modern banks*.

Johri, A., Sayal, A., Chong, K. M., Khoja, M., Jha, J., & Tyagi, N. (2026). Enhancing audit quality and reducing costs: The impact of AI in banking and financial services. *Frontiers in Artificial Intelligence, 8*, 1718854.

Jun, C., & Paulson, N. (2025). *Artificial intelligence and data-driven banking transformation: A review of efficiency and innovations* (SSRN No. 5413924). SSRN.

Kalaiselvi, R. (2026). Explainable AI models for cloud-based fraud detection risk assessment and secure financial decision intelligence. *International Journal of Technology, Management and Humanities, 12*(2), 35–46.

Kumar, S. A. (2026). Artificial intelligence-driven banking and enterprise innovation through secure computing and data-centric systems. *International Journal of Science, Research and Technology, 9*(3), 844–854.

Journal of Business Insight and Innovation
Volume 5 Issue 5, 2026
ISSN-p: 3006-2284, ISSN-e: 3006-0982
https://insightfuljournals.com/

Machireddy, J. R., Rachakatla, S. K., & Ravichandran, P. (2021). AI-driven business analytics for financial forecasting: Integrating data warehousing with predictive models. *Journal of Machine Learning in Pharmaceutical Research, 1*(2), 1–24.

Malik, M. (2025). *Technological and data-driven management: Artificial intelligence in complex decision-making*. In *Navigation complexity: Multidisciplinary approaches to management*.

Mehmood, F. (2024). AI-driven decision support systems in financial risk management: A comparative study. *Multidisciplinary Research in Computing Information Systems, 4*(2), 98–109.

Milton, T. (2025). Artificial intelligence in economic and financial decision-making: A comprehensive review. *Stout in Economics, Finance and Accounting, 1*(1), 1–17.

Nuritdinovich, M. A., Bokhodirovna, K. M., Kavitha, V. O., & Ugli, S. A. O. (2025, June). Advanced AI algorithms in accounting: Redefining accuracy and speed in financial auditing. In *AIP Conference Proceedings* (Vol. 3306, No. 1, Article 050008). AIP Publishing.

Oko-Odion, C. (2025). AI-driven risk assessment models for financial markets: Enhancing predictive accuracy and fraud detection. *International Journal of Computer Applications Technology and Research, 14*(4), 80–96.

Olowe, K. J., Edoh, N. L., Zouo, S. J. C., & Olamijuwon, J. (2024). Review of predictive modeling and machine learning applications in financial service analysis. *Computer Science & IT Research Journal, 5*(11), 2609–2626.

Paleti, S. (2022). Fusion bank: Integrating AI-driven financial innovations with risk-aware data engineering in modern banking. *Decision Making, 2326*, 9865.

Paleti, S. (2024). *Transforming financial risk management with AI and data engineering in the modern banking sector*. American Journal of Analytics and Artificial Intelligence.

Peace, N., Luke, I., & Smith, R. (n.d.). *AI-powered financial forecasting and predictive analytics in digital banking environments*.

Pillai, V. (2023). Integrating AI-driven techniques in big data analytics: Enhancing decision-making in financial markets. *International Journal of Engineering and Computer Science, 12*(7), 10–18535.

Purwar, M., Deka, U., & Raj, H. (2024, November). Data-driven insights: Leveraging analytics for predictive modeling in finance. In *2024 4th International Conference on Technological Advancements in Computational Sciences (ICTACS)* (pp. 687–693). IEEE.

Sundar, J. S., Chandra, S., Saini, R. K., Jha, S., & Katta, S. K. (2025, July). AI-powered predictive analytics for financial risk management in banking and fintech. In *2025 IEEE 4th World Conference on Applied Intelligence and Computing (AIC)* (pp. 174–178). IEEE.

Yanney, A. A. S. (2025). Redefining corporate financial governance through AI-powered predictive models for global business risk management. *International Journal of Research Publication and Reviews, 2*(6), 25–49.