International Journal of Computer Applications Technology and Research
Volume 13–Issue 12, 298 – 308, 2024, ISSN:-2319–8656
DOI:10.7753/IJCATR1312.1023

Predictive Anomaly Detection for Real-Time Virtual Account Transactions and API-Driven Payment Processing Across Enterprise Banking Platforms

Racheal Kikachukwu Ogan
Stanbic IBTC
Lagos, Nigeria

Abstract: The migration of enterprise banking toward instant payments and virtual-account structures has changed anomaly detection from retrospective transaction review into a real-time analytical problem. Virtual accounts can generate high-frequency payment streams across corporate collections, reconciliation processes, customer accounts, and API-connected banking services, creating transaction patterns that static rules may inadequately distinguish from operational or behavioral abnormalities. This study develops a predictive anomaly-detection framework for real-time virtual-account transactions and API-driven payment processing. Rather than evaluating transactions independently, the framework models temporal and contextual behavior using transaction velocity, amount deviation, beneficiary patterns, account activity, payment sequences, processing status, API latency, response codes, retry frequency, and failed-payment clusters. Historical and rolling behavioral baselines establish expected activity at account, transaction, and API levels, while machine-learning and anomaly-detection models generate dynamic anomaly scores for incoming events. Temporal validation, threshold calibration, and false-positive analysis evaluate detection reliability under changing transaction conditions. The resulting intelligence differentiates account-level anomalies from payment-processing irregularities, supporting prioritized investigation, adaptive controls, faster exception management, and resilient real-time banking operations.
Keywords: virtual accounts; predictive anomaly detection; real-time payments; transaction intelligence; payment APIs; banking platforms

1. INTRODUCTION

1.1 Real-Time Banking, Virtual Accounts, and API-Driven Payments

Enterprise banking has increasingly shifted toward interconnected payment environments combining virtual accounts, instant-payment infrastructure, open APIs, payment gateways, treasury platforms, enterprise applications, and automated reconciliation systems [1]. Virtual accounts provide identifiable transaction references that enable incoming payments to be associated with customers, invoices, business units, or commercial activities without maintaining separate physical bank accounts for every relationship [2]. This granularity improves automated allocation and reconciliation while generating detailed behavioural histories for analytical monitoring. Simultaneously, API connectivity enables payment instructions, account information, authorization responses, status updates, and reconciliation events to move rapidly between banks and enterprise systems [3]. A single payment may consequently generate multiple interconnected events across authentication, initiation, authorization, routing, settlement, notification, and reconciliation stages. As transaction volumes expand, financial institutions must process increasingly high-velocity streams containing behavioural and technical information [4]. This environment creates opportunities for automation while simultaneously increasing the complexity of distinguishing normal transaction variation from emerging abnormalities in real time [5].

1.2 The Real-Time Anomaly Detection Problem

Real-time payment environments create an anomaly-detection problem because transaction behaviour can change more rapidly than static monitoring rules can be redesigned [6]. Fixed thresholds may identify transactions exceeding predetermined values or frequencies, but they can struggle to distinguish genuine risk from legitimate changes in customer behaviour [7]. An unusual but legitimate transaction may reflect seasonal purchasing, a new supplier, increased business activity, or an atypical payment destination. Account-behaviour anomalies instead represent departures from established transaction patterns, while payment-processing anomalies can arise through unusual authorization, routing, settlement, or reconciliation events [8]. API-related irregularities may include abnormal request frequency, repeated retries, unexpected response sequences, latency deterioration, or inconsistent transaction states. Potentially high-risk patterns can combine several behavioural and processing signals that individually appear inconclusive [9]. Excessive sensitivity creates false positives, increasing investigation workloads and potentially interrupting legitimate payments. Conversely, insensitive controls may allow consequential anomalies to proceed undetected. Effective detection must therefore balance early identification with operational continuity and proportionate intervention [1].

1.3 Objective, Scope, and Contribution

This study develops a predictive anomaly-detection framework combining virtual-account behavioural patterns with API-level payment-processing signals to identify abnormal transactions in real time [7]. Rather than treating transaction behaviour and technical processing telemetry as

International Journal of Computer Applications Technology and Research
Volume 13–Issue 12, 298 – 308, 2024, ISSN:-2319–8656
DOI:10.7753/IJCATR1312.1023

independent monitoring environments, the framework integrates both within a common analytical architecture. Virtual-account histories provide information on transaction amount, frequency, counterparties, timing, velocity, and behavioural deviation, while API telemetry captures processing sequences, response conditions, retries, latency, authentication events, and transaction-state irregularities [3]. Dynamic baselines establish expected behaviour for individual accounts and relevant operating contexts, enabling deviations to be evaluated relative to historical patterns rather than universal thresholds [9]. These signals feed anomaly intelligence and predictive risk scoring that support investigation and proportionate intervention. Confirmed outcomes subsequently return to the analytical environment for monitoring and model refinement [5].

2. VIRTUAL ACCOUNT TRANSACTION ARCHITECTURE AND ANOMALY DYNAMICS

2.1 End-to-End Virtual Account Payment Architecture

Virtual-account payment environments connect originating customers or businesses with virtual-account structures, banking platforms, payment APIs, core-processing systems, settlement mechanisms, reconciliation services, and enterprise applications [8]. A transaction begins when a customer initiates a payment associated with an identifiable virtual account or payment reference. Authentication establishes whether the initiating party or application is permitted to submit the instruction, after which routing determines the appropriate banking service, processor, or payment rail [6]. Authorization evaluates whether processing conditions are satisfied before the transaction is posted to the relevant account records. Subsequent settlement transfers the financial obligation between participating institutions, while notifications communicate transaction status to connected applications [11]. Reconciliation subsequently matches banking records with enterprise invoices, orders, customer accounts, or accounting entries. Some interactions are synchronous, returning immediate API responses within the requesting session, whereas asynchronous events, including webhooks, settlement confirmations, reversals, and reconciliation updates, may arrive later [9]. Consequently, one payment generates a temporally distributed sequence of interconnected transaction and processing events.

2.2 Transaction and Account-Behaviour Anomalies

Virtual-account activity can exhibit abnormalities across transaction value, frequency, timing, counterparties, direction, and sequential behaviour [12]. Unusual transaction amounts arise when payment values depart materially from an account's established distribution or relevant contextual baseline. Abnormal velocity occurs when transaction frequency increases or decreases unexpectedly within short observation windows, while timing anomalies reflect activity during periods inconsistent with historical patterns [7]. Beneficiary deviation can identify payments directed toward previously unseen or infrequently used counterparties. Sudden changes in account activity may include rapid increases in transaction volume, altered inflow-to-outflow relationships, or unexpected movement following prolonged inactivity [14]. Repeated transactions can represent legitimate recurring payments but may also indicate duplication or unusual behavioural concentration. Dormant-account activation becomes informative when renewed activity differs substantially from the account's previous operating characteristics [10]. Abnormal transaction sequences can additionally reveal combinations of events that appear ordinary individually but unusual collectively. Importantly, these deviations should be interpreted as risk signals requiring contextual evaluation rather than direct evidence of fraudulent activity [6]. Legitimate business expansion, seasonal activity, or changing payment relationships may produce similar deviations.

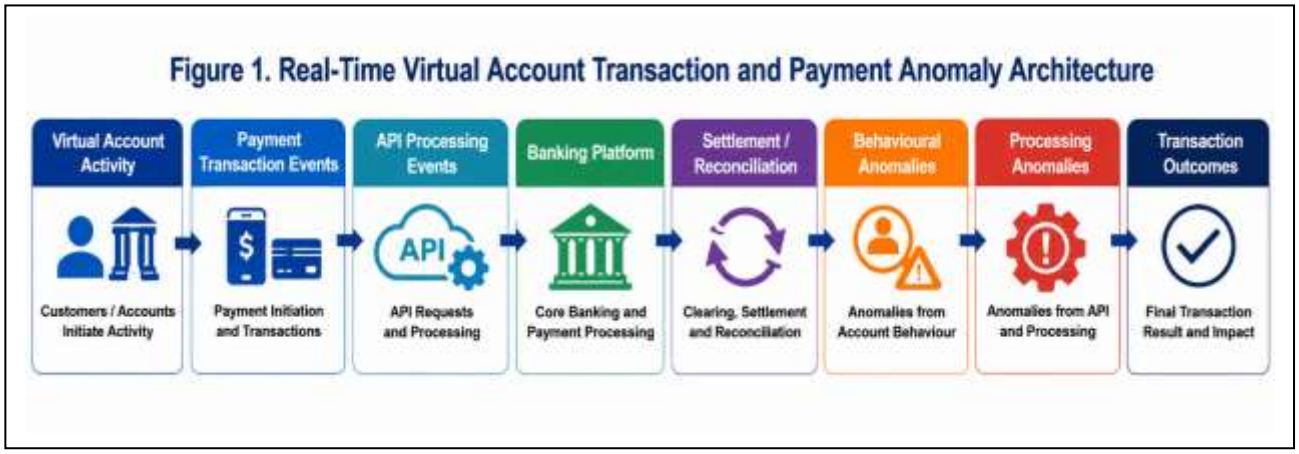

2.3 API and Payment-Processing Anomalies

Behavioural anomalies should be analytically distinguished from abnormalities originating within the technical payment-processing environment [13]. Processing anomalies can include unusual API latency, repeated requests, unexpected response codes, escalating retries, authorization failures, duplicate processing, incomplete transaction states, delayed callbacks, reconciliation mismatches, and inconsistencies between transaction-status records [8]. Elevated latency may indicate dependency deterioration even when a payment eventually completes successfully, while repeated API

requests can arise from timeout handling, client retries, or communication instability. Duplicate processing may occur when retry logic creates multiple instructions without effective idempotency controls [11]. Similarly, delayed callbacks can cause enterprise applications to retain outdated transaction states despite successful downstream processing. Reconciliation mismatches arise when banking, settlement, and enterprise records do not resolve to a consistent outcome [14]. Without analytical separation, these technical abnormalities may resemble suspicious customer behaviour. For example, repeated payment attempts could reflect deliberate account activity or automated retries following processing failure [9]. Combining behavioural and processing evidence therefore requires preserving their distinct origins before assessing their joint contribution to transaction risk.

3. TRANSACTION DATA INTEGRATION AND DYNAMIC ANOMALY FEATURE ENGINEERING

3.1 Multi-Source Transaction and API Data Acquisition

Real-time anomaly detection requires an integrated dataset combining financial behaviour with the technical events generated during payment processing [16]. Transaction sources should include virtual-account transactions, timestamps, amounts, balances, beneficiaries, payment channels, transaction statuses, authorization events, reconciliation records, account histories, and confirmed incident outcomes. Processing telemetry should incorporate API requests and responses, latency, response codes, retries, authentication events, and payment failures [13]. Reliable linkage depends on identifiers preserved across the payment lifecycle. Transaction IDs distinguish individual financial events, while virtual-account identifiers connect transactions with historical account behaviour. Payment references associate instructions with corresponding commercial obligations or reconciliation records [19]. API correlation IDs trace related requests and responses across distributed banking services, supporting identification of processing abnormalities that emerge outside the originating application. Precise timestamps establish when initiation, authorization, posting, callback, settlement, and reconciliation events occurred [15]. Together, these identifiers permit behavioural and processing observations to be analyzed at transaction, account, interface, and temporal levels while preserving their original operational context and supporting subsequent anomaly attribution [20].

3.2 Transaction Sequence Reconstruction and Temporal Integrity

Individual records should be reconstructed into complete payment-event sequences before predictive features are generated [14]. Each journey should connect transaction initiation, authentication, API submission, authorization, routing, posting, status notification, settlement, and reconciliation according to transaction identifiers and timestamps. Chronological ordering is essential because the analytical meaning of an event depends partly on what occurred beforehand [18]. Duplicate records should be distinguished from genuine repeated payment attempts, while missing events require separation from confirmed processing failures. Delayed messages and asynchronous callbacks can arrive after downstream events have already occurred, requiring event-time rather than database-arrival order to determine the actual processing sequence [12]. Temporal integrity also requires an explicit information cutoff. A prediction generated at time $t$should contain only transaction, account, and processing information observable at or before t[20]. Subsequent settlement results, confirmed incidents, chargebacks, or investigation outcomes must not enter historical predictors. This restriction prevents future-information leakage and ensures that retrospective model evaluation reproduces the information available when real-time decisions would actually have been made [17].

3.3 Behavioral Feature Engineering

Behavioural feature engineering should transform raw virtual-account activity into indicators representing how current transactions differ from evolving account-specific expectations [19]. Transaction velocity can measure the number or value of payments occurring within defined short-term windows, while amount deviation quantifies how far a transaction departs from the account's historical value distribution. Balance change captures unusual movements in available funds surrounding payment activity [13]. Beneficiary novelty indicates whether a counterparty is new, infrequent, or inconsistent with established relationships, whereas time-of-day deviation measures departure from typical transaction timing. Transaction-frequency change captures acceleration or contraction in activity relative to preceding periods [16]. Inflow/outflow imbalance can identify unusual directional movement of funds, and rolling transaction volume summarizes recent activity across defined windows. Repeated-payment frequency measures unusually similar payments occurring within compressed periods, while account inactivity duration provides context for renewed activity following dormancy [18]. Transaction-sequence deviation can compare current event patterns with historically observed payment sequences. These indicators should not rely exclusively on fixed historical averages. Rolling windows allow expected behaviour to adapt as legitimate account activity changes, while lagged features preserve preceding conditions without introducing future observations [14]. Multiple windows can distinguish immediate abnormalities from persistent behavioural shifts. Interaction variables can additionally combine amount deviation with beneficiary novelty, velocity with dormant-account activation, or balance change with unusual transaction timing, capturing compound anomalies that isolated indicators may not adequately represent [20].

3.4 API and Processing-Risk Features

Processing-risk features should represent abnormal conditions within the technical infrastructure supporting individual payment journeys [12]. API latency deviation measures departure from contextually expected response times, while

International Journal of Computer Applications Technology and Research
Volume 13–Issue 12, 298 – 308, 2024, ISSN:-2319–8656
DOI:10.7753/IJCATR1312.1023

error-rate acceleration captures whether unsuccessful responses are increasing rather than merely remaining elevated. Retry intensity and timeout frequency indicate communication instability, and authentication-failure rate measures repeated unsuccessful identity or access validation [17]. Duplicate-request frequency can identify retry defects or ineffective idempotency controls. Status-transition irregularity captures unexpected movement, repetition, or absence among transaction states, while callback delay measures abnormal elapsed time before asynchronous confirmation [15]. Processing-duration deviation identifies transactions remaining unusually long within particular stages, and reconciliation-mismatch frequency measures persistent inconsistencies between banking, settlement, and enterprise records. Interaction features should connect these technical indicators with behavioural conditions [19]. For example, elevated transaction velocity combined with rising API latency may represent a different risk mechanism from either condition independently. Similarly, repeated payments accompanied by duplicate API requests may indicate processing instability rather than abnormal customer behaviour, enabling more precise separation of behavioural and technical anomaly sources [13].

Table 1. Transaction and API Data Sources, Engineered Anomaly Indicators, and Analytical Significance

Data Source	Raw Variable/ Event	Engineered Indicator	Anomaly Dimension	Analytical Level	Risk Interpretation
Virtual-account records	Amount, balance, timestamp	Amount deviation; balance change	Behavioural	Transaction/account	Departure from normal account activity
Transaction history	Frequency, timing, value	Velocity; rolling volume	Behavioural	Account	Accelerating or abnormal activity
Beneficiary records	Counterparty identifier	Beneficiary novelty	Behavioural	Transaction	New or unusual payment relationship
Payment records	Payment reference, status	Repeated-payment frequency	Behavioural/processing	Transaction	Possible duplication or abnormal repetition
API telemetry	Response time, request status	Latency deviation; error acceleration	Processing	Interface	Deteriorating technical performance
Retry/timeout logs	Retry and timeout events	Retry intensity; timeout frequency	Processing	Interface	Communication instability
Authentication logs	Authentication outcome	Authentication-failure rate	Security/processing	Transaction/interface	Abnormal access or validation failure
Callback/status logs	State and callback timestamps	Callback delay; transition irregularity	Processing	Transaction journey	Incomplete or inconsistent processing
Reconciliation records	Matched/unmatched records	Mismatch frequency	Processing/outcome	Transaction/account	Inconsistent financial records
Incident outcomes	Confirmed anomaly status	Validated outcome label	Outcome	Transaction/account	Ground truth for model development

4. PREDICTIVE ANOMALY DETECTION FRAMEWORK

4.1 Dynamic Behavioral Baselines

Anomaly detection evaluates each transaction relative to dynamic expectations reflecting the heterogeneous characteristics of virtual accounts and payment environments [23]. The framework establishes account-specific, virtual-account, customer-segment, transaction-type, and payment-channel baselines. Account-specific baselines characterize normal transaction amounts, frequencies, beneficiaries, timing, balances, and transaction sequences, while segment-

level baselines provide contextual comparisons for accounts with limited historical observations [19]. Transaction-type and payment-channel baselines capture systematic differences among transfers, collections, recurring payments, instant payments, and other activities. Current observations are compared with rolling historical distributions and contextually similar transactions [25]. Fixed thresholds inadequately represent these heterogeneous populations because an amount, frequency, or timing pattern considered abnormal for one account may represent routine activity for another. Rolling baselines incorporate changing legitimate behaviour and distinguish persistent behavioural evolution from abrupt deviations. This contextual normalization reduces unnecessary anomaly classification while retaining sensitivity to meaningful account-level changes [21].

4.2 Candidate Anomaly-Detection Models

Model development uses a progressive benchmark ladder that compares detection methods with different assumptions and levels of complexity [20]. M0 applies static rules based on predetermined transaction-value, frequency, velocity, or processing thresholds and represents the conventional monitoring benchmark. M1 uses statistical deviation models that quantify departures from account-specific or contextual distributions [24]. M2 applies Isolation Forest to detect unusual observations through isolation characteristics without requiring confirmed anomaly labels. M3 incorporates Local Outlier Factor or clustering-based detection to identify transactions that differ substantially from comparable behavioural observations [22]. M4 uses autoencoder-based anomaly detection, where reconstruction error measures departure from learned representations of predominantly normal transaction activity. M5 applies supervised machine-learning classification when reliable historical labels distinguish confirmed anomalous events from legitimate transactions [25]. Supervised models learn relationships

4.3 Composite Transaction Anomaly Scoring

The analytical framework combines distinct anomaly dimensions into a composite transaction score [23]:

$$A_i = w_B B_i + w_P P_i + w_C C_i + w_H H_i$$

where A_i denotes the composite anomaly score for transaction i, B_i represents behavioural deviation, P_i represents processing abnormality, C_i captures contextual risk, and H_i represents historical risk. The corresponding weights determine the relative contribution of each component [21]. Weight estimation uses validation performance, confirmed historical outcomes, expert constraints, or optimized combinations of these inputs. Score calibration converts heterogeneous model outputs into comparable scales and interpretable risk levels [25]. Stability analysis across accounts, transaction types, channels, and operating periods identifies whether individual components exert disproportionate influence. Composite scoring therefore preserves the analytical distinction between behavioural and technical abnormalities while integrating them into a unified transaction-risk measure for real-time assessment.

4.4 Temporal Training, Validation, and Threshold Optimization

Model development preserves chronological transaction ordering through sequential training, validation, and untouched out-of-time testing [24]. Earlier observations establish behavioural baselines and model parameters, subsequent observations support hyperparameter and threshold selection, and the final period provides independent performance assessment. Rolling-origin validation evaluates detection stability across successive operating periods as account behaviour, transaction volumes, and payment-processing conditions change [20]. Feature construction,

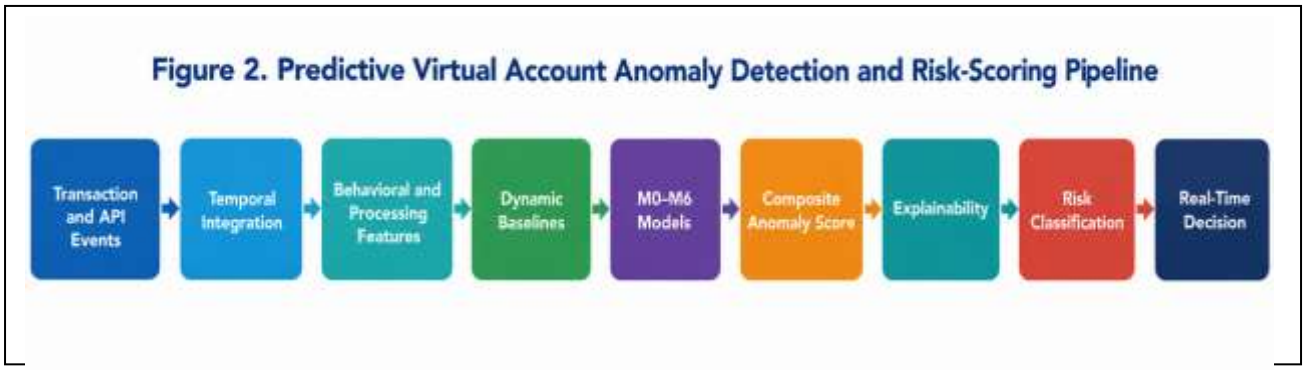

be ıse
un ing

labelled outcomes. M6 combines complementary outputs from multiple detection methods into an ensemble score. Benchmark comparison determines whether added model complexity produces stronger discrimination, stability, and operational usefulness than simpler approaches [19].

development period, preventing future-information leakage. Threshold optimization incorporates both statistical detection performance and operational consequences [25]. Candidate thresholds are evaluated against missed-anomaly costs, false-positive volumes, investigation workload, payment-interruption consequences, transaction exposure, and available intervention capacity. This approach prevents overall accuracy

from dominating model selection and aligns anomaly classification with the economic and operational consequences associated with real-time banking decisions [22].

4.5 Explainable Anomaly Attribution

Explainable attribution identifies the factors responsible for elevated transaction anomaly scores [19]. SHAP values, feature contributions, and interpretable reason indicators separate behavioural abnormalities from technical processing deterioration. A transaction characterized by unusual amount, new beneficiary, and abnormal transaction velocity represents a different anomaly mechanism from one exhibiting normal customer behaviour, excessive API retries, and delayed processing [23]. Local attribution identifies the dominant contributors to individual classifications, while global attribution establishes the variables consistently associated with abnormal transactions across the monitored population [25]. This distinction supports appropriate investigation routing and prevents payment-processing defects from being interpreted automatically as suspicious customer behaviour.

5. REAL-TIME RISK CLASSIFICATION AND PAYMENT INTERVENTION

5.1 Context-Aware Transaction Risk Stratification

Composite anomaly scores require contextual interpretation because statistical abnormality does not necessarily represent equivalent financial or operational risk across transactions [27]. The framework classifies transactions into Normal, Elevated, High, and Critical risk tiers using anomaly magnitude together with transaction amount, account history, beneficiary characteristics, payment channel, behavioural deviation, processing abnormality, and potential financial exposure. Normal transactions remain consistent with established behavioural and processing conditions, whereas Elevated transactions contain moderate deviations requiring increased observation [30]. High-risk transactions combine substantial abnormalities with meaningful exposure or contextual concerns. Critical transactions exhibit severe or interacting indicators associated with significant potential financial consequences [25]. Risk stratification therefore avoids treating every statistical outlier identically. An unusually large payment to an established beneficiary through a historically consistent channel differs materially from the same amount involving a new beneficiary, abnormal velocity, and authentication irregularities [32]. Context-aware classification consequently connects analytical abnormality with the significance of the underlying transaction and supports differentiated operational treatment.

5.2 Risk-Based Intervention

Risk-based intervention links each classification with actions proportionate to the source, severity, and potential consequence of the detected anomaly [29]. Normal transactions receive Normal Processing, while Elevated classifications receive Enhanced Monitoring where deviations remain limited or contextually explainable. High-risk behavioural patterns receive Additional Verification or Analyst Review, particularly where unusual amounts, new beneficiaries, transaction acceleration, or account-history deviations occur [26]. Critical behavioural anomalies receive Temporary Hold or Escalation when financial exposure and supporting risk indicators justify stronger control. Processing abnormalities follow a distinct intervention pathway. Excessive API retries, latency deterioration, callback delays, duplicate requests, or reconciliation inconsistencies trigger Processing Investigation rather than automatic adverse treatment of the account holder [31]. Technical teams examine payment infrastructure, authentication services, APIs, gateways, or downstream dependencies associated with the abnormal processing condition. This distinction prevents system-generated irregularities from being incorrectly interpreted as suspicious customer conduct [24]. Intervention therefore reflects both risk magnitude and anomaly origin, aligning customer-focused controls with behavioural evidence and technical remediation with processing evidence while preserving legitimate payment continuity wherever the available evidence supports normal processing.

5.3 False-Positive Management and Human Review

False-positive management balances detection sensitivity against the operational requirement to process legitimate transactions without unnecessary interruption [28]. Alert prioritization ranks cases according to anomaly score, financial exposure, contextual evidence, explanation strength, and intervention urgency rather than presenting analysts with undifferentiated alerts. Available analyst capacity also influences thresholds because excessive alert volumes reduce investigation depth and delay attention to consequential cases [32]. Explainable indicators provide reviewers with the principal behavioural and processing factors underlying each classification, supporting contextual assessment against account history, beneficiary relationships, payment characteristics, and technical conditions. Human review remains particularly important for High and Critical cases where automated evidence contains uncertainty or conflicting signals [25]. Authorized override mechanisms permit reviewers to modify automated classifications when documented contextual evidence supports a different conclusion. Review records preserve the original score, explanatory factors, analyst assessment, override rationale, and final action [30]. This documentation supports subsequent model evaluation and distinguishes genuine detection failures from legitimate exceptions, strengthening anomaly detection while maintaining transaction continuity.

6. PERFORMANCE EVALUATION AND MODEL VALIDATION

6.1 Anomaly-Detection Performance

For labelled transaction outcomes, detection performance is evaluated through precision, recall, specificity, and F1-score [34]. Precision measures the proportion of flagged transactions confirmed as consequential anomalies:

$$Precision = \frac{TP}{TP + FP}$$

while recall measures the proportion of confirmed anomalies successfully detected:

$$Recall = \frac{TP}{TP + FN}$$

Specificity evaluates correct recognition of legitimate transactions:

$$Specificity = \frac{TN}{TN + FP}$$

The F1-score balances precision and recall through their harmonic mean [31]:

$$F1 = 2\left(\frac{Precision \times Recall}{Precision + Recall}\right)$$

ROC-AUC evaluates discrimination across classification thresholds, whereas PR-AUC emphasizes the precision-recall relationship and provides particular value when consequential anomalies represent a small proportion of transaction activity [38]. False-positive and false-negative rates are additionally reported as:

$$FPR = \frac{FP}{FP + TN}, \; FNR = \frac{FN}{FN + TP}$$

Detection lead time measures the interval between the initial model warning and confirmed anomaly occurrence [35], capturing whether detection occurs early enough to support meaningful operational intervention.

6.2 Transaction and Operational Performance

Statistical discrimination alone does not establish operational value. Evaluation therefore examines anomaly capture, investigation precision, transaction completion, payment-exception resolution, reconciliation accuracy, alert volumes, investigation time, and processing reliability [32]. Anomaly capture is expressed as:

$$ACR = \frac{\text{Detected Confirmed Anomalies}}{\text{Total Confirmed Anomalies}}$$

Investigation precision measures productive use of analyst capacity:

$$IP = \frac{\text{Confirmed Anomalies Investigated}}{\text{Total Alerts Investigated}}$$

Transaction completion and reconciliation accuracy provide complementary measures of payment continuity and processing integrity [39]:

$$TCR = \frac{\text{Successfully Completed Transactions}}{\text{Total Initiated Transactions}}$$

$$RA = \frac{\text{Correctly Reconciled Transactions}}{\text{Transactions Requiring Reconciliation}}$$

Pre-deployment and post-deployment comparisons identify changes in these indicators relative to existing rule-based or operational benchmarks [36]. Percentage performance improvement is measured as:

$$\Delta P(\%) = \frac{P_{post} - P_{pre}}{P_{pre}} \times 100$$

For metrics where lower values represent improvement, including investigation time and false-alert volume, the direction of interpretation is reversed. Evaluation consequently determines whether improved anomaly detection coincides with stronger processing reliability rather than merely producing additional alerts [30].

6.3 Stability, Drift, and Recalibration

Performance stability is examined across changing transaction volumes, customer behaviour, payment channels, API versions, seasonal conditions, and emerging transaction patterns [37]. Population Stability Index provides one measure of distributional change:

$$PSI = \sum_{j=1}^{K} (A_j - E_j) \ln\left(\frac{A_j}{E_j}\right)$$

where A_j represents the observed proportion in current period bin j, and E_j represents the corresponding reference-period proportion. Changes in feature distributions are assessed alongside anomaly-score distributions, calibration, false-positive rates, recall, and operational outcomes [33]. Material deterioration initiates threshold recalibration, model retraining, or comparison with challenger specifications. Challenger models provide evidence on whether observed degradation reflects temporary behavioural change or reduced suitability of the incumbent specification [40]. Periodic out-of-time validation preserves temporal realism by testing performance on observations excluded from model development. Recalibration therefore forms part of continuous model control rather than a response restricted to complete model failure [35].

Table 2. Multi-Level Evaluation of Anomaly Detection and Payment-Processing Performance

Evaluation Level	Metric	Measurement	Benchmark	Interpretation	Operational Significance
Detection	Precision, Recall, F1	Classification outcomes	Existing rules	Detection quality	Reliable anomaly identification
Imbalance	PR-AUC	Precision-recall curve	Baseline model	Rare-event discrimin	Consequential anomaly

Evaluation Level	Metric	Measurement	Benchmark	Interpretation	Operational Significance
				ation	capture
Error	FPR, FNR	Classification errors	Historical rates	Error trade-off	Interruption versus missed risk
Timeliness	Detection lead time	Warning-to-event interval	Existing detection process	Warning speed	Intervention opportunity
Transaction	TCR	Completed/initiated transactions	Pre-deployment rate	Payment continuity	Customer and payment reliability
Investigation	Investigation precision	Confirmed/investigated alerts	Existing workflow	Alert usefulness	Analyst efficiency
Reconciliation	RA	Correctly reconciled/required	Historical accuracy	Processing integrity	Financial control
Stability	PSI and performance drift	Distribution/performance change	Development baseline	Model stability	Recalibration requirement

7. GOVERNANCE, EXPLAINABILITY, AND ENTERPRISE BANKING RESILIENCE

7.1 Explainability and Decision Auditability

Explainability formed an essential control layer because transaction-level intervention required more than an elevated anomaly score. Each decision incorporated reason indicators identifying the behavioural or processing variables contributing most strongly to the classification [39]. Model versioning preserved the exact specification, parameters, thresholds, and feature configuration applied when each transaction was evaluated. Decision logs recorded transaction identifiers, anomaly scores, risk tiers, timestamps, explanations, and resulting actions [42]. Intervention records additionally documented verification requests, processing investigations, temporary holds, escalations, and subsequent resolution. Where analysts overrode automated classifications, the original recommendation, revised decision, reviewer identity, and justification were retained [37]. Outcome histories connected these records with confirmed anomalies, false positives, processing incidents, and legitimate transactions. This structure established traceability from raw transaction evidence through feature construction, anomaly scoring, explanation, human review, intervention, and verified outcome, supporting reconstruction and audit of individual decisions [44].

7.2 Responsible Automated Decision-Making

Responsible deployment required anomaly detection to operate within explicit privacy, security, fairness, proportionality, and human-oversight controls [40]. Data minimization restricted analytical inputs to information relevant to transaction-risk assessment, while access and security controls protected behavioural histories and payment-processing telemetry. Fairness required particular attention because legitimate account populations exhibited materially different transaction frequencies, values, counterparties, channels, and operating schedules [43]. Behaviour that differed from a population-wide pattern was therefore not automatically interpreted as elevated risk. Dynamic account-specific and context-specific baselines reduced inappropriate classification arising from legitimate heterogeneity. Intervention severity was also proportionate to available evidence and potential exposure [38]. Elevated statistical abnormality supported additional monitoring without necessarily interrupting payment activity, while consequential or interacting indicators justified stronger review. Human oversight remained central to High and Critical classifications, especially where behavioural and technical evidence conflicted [45]. Automated scoring consequently supported accountable decision-making rather than replacing contextual professional judgment.

7.3 Closed-Loop Anomaly Learning

Closed-loop learning incorporated confirmed anomaly outcomes, false positives, analyst decisions, processing incidents, transaction completions, and intervention results into subsequent analytical assessment [41]. Newly observed evidence updated behavioural baselines, processing expectations, risk thresholds, and model parameters. Repeated deviations that became established account behaviour were distinguished from temporary abnormalities associated with isolated transactions, seasonal activity, technical incidents, or unusual operating conditions [44]. This distinction reduced persistent false alerts when legitimate behaviour changed while preserving sensitivity to genuinely abnormal events. Differences between predicted and verified outcomes also provided direct evidence for recalibration and challenger-model evaluation [39].

7.4 Enterprise Payment Resilience

The integrated framework connected earlier anomaly identification with stronger enterprise payment resilience [42]. Behavioural and processing intelligence supported reduced

International Journal of Computer Applications Technology and Research
Volume 13–Issue 12, 298 – 308, 2024, ISSN:-2319–8656
DOI:10.7753/IJCATR1312.1023

unnecessary payment disruption, stronger transaction integrity, faster investigation, improved exception resolution, and more reliable processing continuity. Combining real-time detection with contextual classification, proportionate intervention, verification, and continuous learning enabled anomaly-management capacity to scale with transaction volume without relying exclusively on static rules or manual transaction screening [45].

distinguishing infrastructure abnormalities from customer-originated events. Fraud and risk teams gained behavioural explanations supporting investigation prioritization, and model validators obtained traceable evidence linking inputs, scores, interventions, and outcomes [43]. Compliance functions benefited from documented decision histories and human-review controls. This shared intelligence reduced analytical separation between functions that traditionally examined transaction behaviour, processing reliability, and

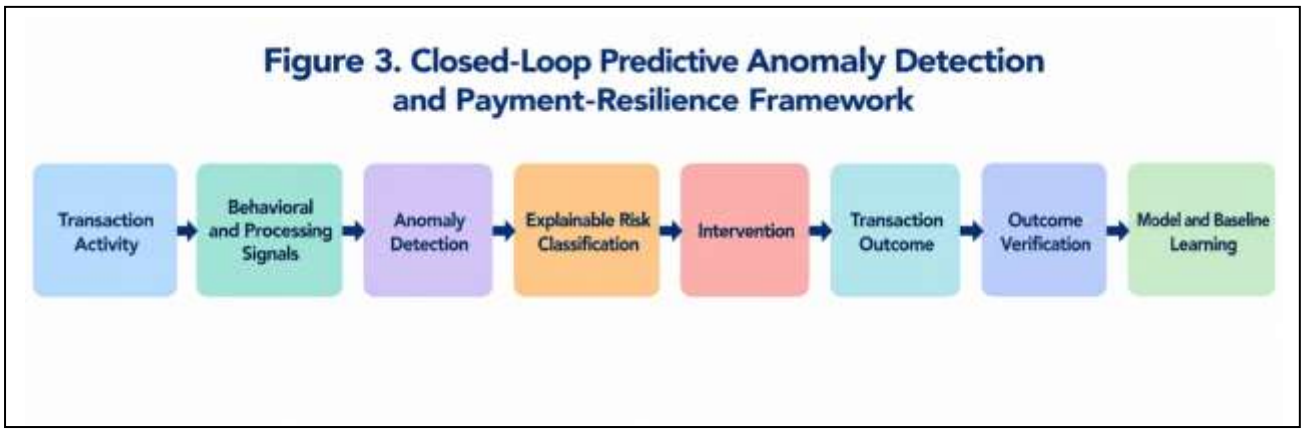

financial risk independently [37]. Operational value therefore arose from coordinated interpretation of transaction and technical evidence rather than isolated optimization of individual monitoring systems.

8. DISCUSSION AND RESEARCH IMPLICATIONS

8.1 Integrated Interpretation of Transaction and Processing Anomalies

The integrated framework demonstrated that virtual-account anomaly detection extended beyond identifying statistically unusual transactions [38]. Effective risk interpretation depended on establishing what was unusual, the baseline against which the deviation occurred, the source of the abnormality, and its potential transaction consequence. Account-behaviour anomalies reflected departures involving amounts, beneficiaries, velocity, timing, frequency, or transaction sequences, whereas processing anomalies originated from latency, retries, authentication failures, callbacks, status transitions, or reconciliation conditions [41]. Separating these dimensions prevented technical irregularities from being incorrectly attributed to customer behaviour. Their integration nevertheless remained important because behavioural and processing abnormalities sometimes occurred simultaneously and jointly increased transaction risk [45]. Anomaly intelligence therefore represented contextual interpretation of multiple signals rather than mechanical identification of statistical outliers.

8.2 Managerial and Technology Implications

The framework established a shared analytical environment for banking technology teams, payment operations, fraud and risk functions, API engineers, transaction-monitoring teams, model validators, compliance personnel, and enterprise architects [40]. Payment operations obtained transaction-level risk indicators, while API teams received evidence

8.3 Limitations and Future Research

Limitations remained around sparse anomaly labels, evolving transaction behaviour, platform heterogeneity, and incomplete attribution of observed abnormalities [44]. Further research could examine sequence learning, graph-based transaction relationships, adaptive anomaly detection, federated learning, causal anomaly attribution, privacy-preserving analytics, real-time model adaptation, and cross-platform transferability [41]. These extensions could strengthen discrimination between persistent behavioural change, interconnected transaction risk, and transient processing abnormalities.

9. CONCLUSION

The study established an integrated relationship among real-time transaction intelligence, dynamic behavioral baselines, processing analytics, predictive anomaly detection, explainable risk classification, proportionate intervention, continuous learning, and payment resilience. Effective anomaly detection extended beyond identifying transactions that departed from statistical norms. It distinguished legitimate behavioral variation from meaningful account-level anomalies and technology-driven payment irregularities while providing interpretable evidence for operational decisions. Integrating virtual-account behavioral signals with API processing telemetry enabled account activity and technical processing conditions to be evaluated within the same analytical environment. This integration supported earlier anomaly identification, more precise investigation, lower false-positive burdens, stronger transaction integrity, and

improved payment continuity. The framework therefore represented a progression from static transaction rules and retrospective exception management toward adaptive, context-aware, and explainable real-time payment intelligence. By connecting detection with proportionate intervention and outcome-based learning, the approach strengthened scalable payment resilience across enterprise banking platforms.

10. REFERENCE

1. Kansara M. Cloud migration strategies and challenges in highly regulated and data-intensive industries: A technical perspective. International Journal of Applied Machine Learning and Computational Intelligence. 2021;11(12):78-121.
2. Christopher JS. A Secure API-Enabled Cloud Platform for Healthcare and Financial Systems Leveraging AI and Apache-Based Real-Time Analytics. International Journal of Research and Applied Innovations. 2024 Jul 8;7(4):11063-72.
3. Usha M. Deep Learning Driven Predictive Threat Detection Framework for Secure Financial and Healthcare Cloud Platforms. International Journal of Engineering & Extended Technologies Research (IJEETR). 2024 Aug 14;6(4):8141-52.
4. Burugulla JK. Cloud-Enabled AI Architectures for Personalized Financial Services and Payment Solutions. International Journal of Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering. 2022 Jan 1.
5. Quân TM. A strategic analysis of AI-driven customer relationship management systems in enhancing personalization and retention in financial institutions. Orient Journal of Emerging Paradigms in Artificial Intelligence and Autonomous Systems. 2024 Sep 4;14(9):1-8.
6. Wiśniewski AT. AI-Driven ERP Cybersecurity for Fuel Cell Vehicles: Real-Time Threat Detection in Oracle and SAP Using NLP and U-Net Denoising. International Journal of Advanced Research in Computer Science & Technology (IJARCST). 2024 Jul 10;7(4):10612-6.
7. Jain R. *Innovative new payment Infrastructure using a new business model with a modern mobile computing and cloud platform* (Doctoral dissertation, University of Portsmouth).
8. Zadeh LA, Caron S, Ngadiuba J. Composable Finance Integrating AI, Blockchain, and APIs for Modular Financial Products. Journal of Integrated Research. 2021 Dec 30;2(4):1-3.
9. Solarin A, Chukwunweike J. Dynamic reliability-centered maintenance modeling integrating failure mode analysis and Bayesian decision theoretic approaches. *International Journal of Science and Research Archive*. 2023 Mar;8(1):136. doi:10.30574/ijsra.2023.8.1.0136.
10. Eyubova NF. THE ESSENCE AND MAIN FEATURES OF THE MODERN BANKING SYSTEM. TİKİNTİNİN İQTİSADİYYATI VƏ MENECMENT № 4-2024.:340.
11. Anand L. AI-powered cloud cybersecurity architecture for risk prediction and threat mitigation in healthcare and finance. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM). 2024 Nov 11;7(Special Issue 1):5-12.
12. Oluwatobi Alebiosu. Engineering stakeholder alignment frameworks for resolving land rights, environmental compliance, and project delivery conflicts. Int J Appl Res 2024;10(12):401-412.
    DOI: 10.22271/allresearch.2024.v10.i12e.13917
13. Timilehin O. The Future of Financial Technology: Emerging Trends and Innovations. no. December. 2024.
14. Emmanuel Uzoh. PREDICTIVE OPERATIONS INTELLIGENCE: AN AUTOMATED MACHINELEARNING SYSTEM FOR BOTTLENECK DETECTION AND DYNAMIC WORKSTREAM PRIORITIZATION IN SMALL AND MEDIUM-SIZED ENTERPRISES. INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM). 2023Dec21;07(12):793–806. doi:10.5281/zenodo.22089397
15. Rajesh SC, Jain U. Real-Time Billing Systems for Multi-Tenant SaaS Ecosystems. International Journal of Science and Research Archive. 2024;13(01):1059-80.
16. Egogo-Stanley AO, Ibrahim OM, Akinyemi AD. Assessing flood vulnerability using GIS spatial analytics to inform infrastructure planning, emergency response and community resilience strategies. Int J Sci Res Arch. 2022;7(2):952-969. doi:10.30574/ijsra.2022.7.2.0355.
17. Baladari V. AI-Powered Chatbots in Banking: Developer Best Practices for Enhancing Efficiency and Security. International Journal of Advanced Research in Science Communication and Technology. 2024;4:649-56.
18. Marchand TP. An AI-Enabled Cloud Framework for Healthcare ERP Targeted Advertising with LLM Analytics and Deep Learning–Based Industrial Effluent Prediction. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM). 2022 Dec 25;5(6):7822-9.
19. Oluwatobi Alebiosu. EVALUATING POLICY-TO-PROJECT EXECUTION GAPS IN NATIONAL ENERGY TRANSITION PROGRAMS THROUGH GOVERNANCE MATURITY ANALYTICS. INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM). 2019Dec21;03(12):191–205. doi:10.5281/zenodo.21454333
20. Onwuegbuchi O, Ibiyeye AO, Okolo JN, Adeniji SA. Cybersecurity risks in the fintech ecosystem: Regulatory and technological perspectives. Communication in Physical Sciences. 2023 Sep 19;9(4):9547-67.
21. AKINYELURE FM. Music and Dance Therapy; The Effectiveness of Occupational Therapy Intervention in Preventing Depression amongst Persons living with Alzheimer in Lagos Nigeria. Alzheimer's & Dementia. 2022 Dec;18:e069054.
22. Bakker SA. An AI-Cloud Converged Architecture for Cybersecurity Fraud Analytics and Medical Imaging over 5G with Oracle EBS and Unified Payment Orchestration. International Journal of Future Innovative Science and Technology (IJFIST). 2023 Jul 6;6(4):10985.

23. Adewuyi A, Nwangele CR, Oladuji TJ, Akintobi AO. Advances in Machine Learning for Credit Risk and Underwriting Automation: Emerging Trends in Financial Services. International Journal of Advanced Multidisciplinary Research and Studies. 2023;3(6):1860-77.
24. Oluwatobi Alebiosu. Modeling Contractual Risk Propagation Mechanisms Within Multi-Stakeholder Energy Transition Megaproject Delivery Frameworks. Int J Comput Artif Intell 2021;2(2):103-114. DOI: 10.33545/27076571.2021.v2.i2a.357
25. Tiwari SK. The Future of digital retirement solutions: A Study of sustainability and scalability in financial planning tools. Journal of Computer Science and Technology Studies. 2024 Dec 25;6(5):229-45.
26. Oluwatosin Michael Ibrahim, Andy Osagie Egogo-Stanley, Ayomide D Akinyemi. (2021). LEVERAGING GEOSPATIAL INFORMATION SYSTEMS FOR PREDICTIVE FLOOD MODELING AND EVIDENCE-DRIVEN DISASTER RISK REDUCTION POLICY DEVELOPMENT. International Journal Of Engineering Technology Research & Management (IJETRM), 05(12), 397–415. https://doi.org/10.5281/zenodo.18378803
27. Jameaba MS. Digitalization, emerging technologies, and financial stability: challenges and opportunities for the Indonesian banking sector and beyond. Emerging Technologies, and Financial Stability: Challenges and Opportunities for the Indonesian Banking Sector and Beyond (April 26, 2024). 2024 Apr 26.
28. Joshi S. Gen AI for market risk and credit risk learn agentically powered gen AI; gen AI agentic framework for financial risk management. Gen AI Agentic Framework for Financial Risk Management (December 15, 2024). 2024 Dec 15.
29. Uzoh E. Predictive credit analytics for strengthening SME lending decisions, portfolio quality, regulatory compliance, and sustainable financial institution performance. Int J Foreign Trade Int Bus. 2022;4(2):72-83. doi:10.33545/26633140.2022.v4.i2a.271.
30. Vemireddy S. Modernizing enterprise financial platforms through distributed cloud architectures. International Journal of Science, Research and Technology. 2022;5(2):7420-6.
31. Temitope Iyamu Fadairo. AI powered customer experience in financial services: Balancing personalization, efficiency, and regulatory compliance. Int J Finance Manage Econ 2023;6(2):237-247. DOI: 10.33545/26179210.2023.v6.i2.635
32. Kiran A. Strengthening Digital Banking: An Integrated Cyber Resilience Framework. International Journal of Recent Trends in Technology and Engineering. 2024 Sep 27;3(3):1-4.
33. Oluwatobi Alebiosu. Quantifying regulatory delay risk transmission across utility-scale renewable energy projects in hydrocarbon-dominant economies. Int J Hydropower Civ Eng 2023;4(1):35-47. DOI: 10.22271/27078302.2023.v4.i1a.89
34. Mohammed A. Open Banking and APIs: Research on how open banking frameworks and APIs are reshaping the financial ecosystem. International Journal of Advances in Engineering and Management. 2024 Sep 13;7:770-84.
35. Igwe-Nmaju C. Organizational communication in the age of APIs: integrating data streams across departments for unified messaging and decision-making. International Journal of Research Publication and Reviews. 2024 Dec;5(12):2792-809.
36. Pinyi EO. Evaluating AI-enabled real-time security architectures for critical infrastructure: an empirical multi-domain study. *Int J Sci Eng Appl*. 2024;13(12):104-117. doi:10.7753/IJSEA1312.1015.
37. Varghese J. Intelligent Cloud Native AI Architectures for Secure Enterprise Data Management and Scalable Financial Systems. International Journal of Science, Research and Technology. 2022;5(6):8930-44.
38. Oluwatobi Alebiosu. Architecting Energy Transition Governance Models for Balancing Decarbonization Objectives with Hydrocarbon Asset Optimization Strategies. Int J Res Civ Eng Technol 2020;1(2):08-20. DOI: 10.22271/27078264.2020.v1.i2a.139
39. Silbermann FD. Cloud-native AI and deep learning models for real-time fraud detection and cybersecurity in financial institutions. International Journal of Advanced Research in Computer Science & Technology (IJARCST). 2024 May 22;7(Special Issue 1):17-25.
40. Akinyemi AD, Opejin A, Egogo-Stanley AO, Ibrahim OM. GIS-enabled flood hazard assessment for enhancing early warning systems, land use regulation, and sustainable watershed management. Global J Eng Technol Adv. 2023;17(3):99-118. doi:10.30574/gjeta.2023.17.3.0262.
41. Okafor CM, Dako OF, Osuji VC. Engineering High-Throughput Digital Collections Platforms for Multi-Billion-Dollar Payment Ecosystems. Shodhshauryam, International Scientific Refereed Research Journal. 2021 Jul;4(4):315-35.
42. Magoni D. AI and Machine Learning Powered Secure Enterprise Platforms with Real-Time APIs Financial Forecasting Risk Analytics and Blockchain. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM). 2024 Nov 22;7(6):11622-30.
43. Uzoh E. Leveraging customer behavioral analytics to strengthen brand governance, seller accountability, and service quality across e-commerce marketplaces. Int J Res Marketing Manage Sales. 2024;6(2 Pt D):347-356. doi:10.33545/26633329.2024.v6.i2d.437.
44. Ononiwu M, Azonuche TI, Okoh OF, Enyejo JO. Machine learning approaches for fraud detection and risk assessment in mobile banking applications and fintech solutions. International Journal of Scientific Research in Science, Engineering and Technology. 2023 Jul;10(4):1-0.
45. Kotla MR. Intelligent cloud-native banking: Leveraging machine learning for secure and scalable digital financial services. International Journal of Engineering & Extended Technologies Research (IJEETR). 2022 Dec 15;4(6):5758-61.