International Journal of Scientific Research and Modern Technology (IJSRMT) ijsrmt.com
Volume 5, Issue 10, 2026
DOI: https://doi.org/10.38124/ijsrmt.v5i10.1747

Real-Time Behavioural Fraud Prevention in Financial Institutions Using Adaptive Machine Learning

Farainesu Hove[1]; Mainford Mutandavari[2]

[1]Information Technology Department, School of Information Science and Technology, Harare Institute of Technology, Harare, Zimbabwe
[2]Department of Data Science and Business Systems, SRM Institute of Science and Technology, Kattankulathar, India

Publication Date: 2026/11/18

Abstract

The proliferation of digital financial channels has exponentially increased the attack surface for financial fraud, with global card fraud losses reaching USD 33.45 billion in 2022. Conventional rule-based fraud detection systems are inherently incapable of adapting to fast-changing fraud tactics, leading to high levels of undetected fraud and too many false positives impacting customer experience. This paper proposes a real-time behavioural fraud prevention framework based on Adaptive Machine Learning (AML) and CRISP-DM methodology. In this study, three classification algorithms, Decision Tree (DT), Random Forest (RF) and Extreme Gradient Boosting (XGBoost) are trained, evaluated and compared using an anonymised dataset of 499,999 financial transactions extracted from a Zimbabwean commercial bank on imbalanced and class-balanced datasets. We engineer novel behavioural features from raw transactional attributes, namely errorBalanceOrg, errorBalanceDest and hourOfDay that account for 71% of the feature importance of the Random Forest model. The Decision Tree classifier gives the best Area Under the Curve (AUC) of 0.9633361 on the balanced data set, which is a significant improvement over the estimated rule-based baselines. Current adaptive ML trends including federated learning, graph neural networks, transformer based anomaly detection, and concept drift adaptation are also discussed in this paper. In addition, a cloud-native deployment architecture on Microsoft Azure is proposed. The results demonstrate the viability of adaptive machine learning as a high-performance alternative to static rule-based fraud detection in developing-economy banking environments.

Keywords: *Adaptive Machine Learning, Fraud Detection, CRISP-DM, Real-Time Prevention, Decision Tree, Random Forest, XGBoost, Concept Drift, Behavioural.*

I. INTRODUCTION

Financial fraud is among the most significant and persistent threats to the global financial ecosystem. The exponential growth of digital banking channels (including mobile banking applications, internet banking platforms, real-time payment services, and electronic funds transfer systems) has significantly expanded the attack surface available to malicious actors. According to the Nilson Report (2023), global card fraud losses amounted to USD 33.45 billion in 2022, increasing at a compound annual growth rate of 9.3% over the last five years. In Zimbabwe, the digital financial revolution has been fast-tracked by the use of mobile money platforms (EcoCash, One-Money), interbank electronic transfers (ZIPIT) and the ubiquitous use of VISA and Mastercard debit services. In 2022 alone, the Reserve Bank of Zimbabwe (RBZ) National Payment Systems Report (2022) recorded over 900 million electronic transactions amounting to USD 145 billion in value, creating an unprecedented volume of transactional data from which fraud patterns must be identified in near real time. The conventional architecture to detect fraud in financial institutions is based on static rule-based expert system capturing the domain knowledge in pre-defined thresholds and logic. Rule-based systems are operationally straightforward, but they have crucial structural limitations: they are not able to learn from emerging fraud patterns, they produce high false-positive

Hove, F., & Mutandavari, M. (2026). Real-Time Behavioural Fraud Prevention in Financial Institutions Using Adaptive Machine Learning. *International Journal of Scientific Research and Modern Technology*, *5*(10), 175–184. https://doi.org/10.38124/ijsrmt.v5i10.1747

rates (typically 3-5%), they need continuous manual maintenance by domain experts, and they do not model the complex non-linear relationships inherent in modern fraudulent transaction behaviour [1, 2].

Page 2 of 11 Adaptive Machine Learning (AML) represents a paradigm shift: systems that autonomously learn from streaming transaction data, continuously adapt to evolving fraud tactics and make high confidence real-time prevention decisions before fraudulent transactions are authorised [3, 4]. Global financial institutions have deployed supervised classification algorithms, unsupervised anomaly detectors and adaptive drift correcting models in production fraud prevention pipelines, and have demonstrated up to 50% reduction in fraud losses compared to rule-based baselines. The paper makes four main contributions to the literature. First, it presents the first empirically validated comparative study of Decision Tree, Random Forest and XGBoost classifiers on financial transaction data from a Zimbabwean commercial bank. Second, it demonstrates the impact of engineered balance-error features on model discriminatory performance. Third, it synthesises current global adaptive ML trends and evaluates their applicability to developing-economy financial environments. Finally, it proposes a cloud-native real-time deployment architecture suitable for integration with existing banking IT infrastructure.

➢ *Background to the Study*

Traditionally, banking fraud mitigation was based on manual verification and slow transaction settlement procedures, giving staff time to review suspicious transactions before they were completed. This time lag does not exist in modern digital banking environments: transactions are processed instantly, and once authentication credentials have been compromised, fraudulent transfers can be completed within seconds. Financial institutions solve this problem mainly with rule-based monitoring systems that compare transactions against predefined conditions — abnormal transaction amounts, unusual geographic locations, blacklisted accounts. Rule-based approaches are effective against known fraud scenarios, but rely entirely on prior knowledge of fraud patterns and can easily be circumvented when fraudsters adapt their strategies. This results in a double failure mode: undetected fraud events and a high volume of false alerts that require expensive manual investigation. Customers have persistent behavioural traits such as transaction frequency, ranges of expenditure, distributions of time-of-day activity, device usage profiles and geographic footprints. Continuous analysis of such behavioural signals enables dynamic, customer-specific fraud detection capabilities that go beyond generic threshold-based rules [5, 6]. Behavioural modelling combined with adaptive machine learning allows financial systems to learn changing transaction behaviour and assess risk at the time of transaction authorisation rather than in post-processing.

➢ *Research Objectives*

This study is guided by the following specific objectives:

- Extract and construct behavioural transaction features — including transaction frequency, spending patterns, device usage, location consistency, and transaction timing — from multi-channel banking transaction data.
- Develop an adaptive machine learning risk-scoring model capable of continuously learning evolving user behaviour and detecting anomalous transaction patterns.
- Implement a real-time transaction decision engine that dynamically approves or blocks transactions based on calculated behavioural risk thresholds.
- Evaluate model performance using Precision, Recall, F1-Score, AUC-ROC, and False Positive Rate in comparison with traditional rule-based detection systems.

➢ *Research Questions*

The central research question is: To what extent can a real-time adaptive behavioural fraud prevention model reduce fraudulent digital banking transactions compared to traditional rule-based detection systems? This overarching question is operationalised through five sub-questions addressing: (i) behavioural modelling accuracy; (ii) adaptive ML detection performance; (iii) real-time prevention capability; (iv) comparative performance against rule-based baselines; and (v) practical fraud reduction effectiveness.

II. RELATED WORK

We review related work in four topics: machine learning for financial fraud detection, real-time detection architectures, modelling behavioural transactions, and adaptive learning under concept drift. The synthesis reveals a persistent gap in the current systems: their role as post-authorisation detection mechanisms rather than pre-authorisation prevention systems.

➢ *Machine Learning in Financial Fraud Detection*

Early fraud detection systems were based on rule-based monitoring techniques that analysed transactions according to predefined thresholds like transaction amount limits, geographic inconsistencies, and blacklisted accounts [8]. Statistical fraud detection is described by Bolton and Hand [12] as the detection of anomalous behaviour in financial data relative to the expected patterns.

These systems work well for known fraud scenarios, but they cannot detect new strategies which have never been seen before, because they are based on manually defined rules that require continuous maintenance by experts. Bhattacharyya et al. [7] showed that machine learning classifiers such as Random Forest and Support Vector Machines significantly outperform rule-based systems in the detection of fraudulent credit card transactions with accuracy improvements of 15–25% on

benchmark datasets. Carcillo et al. [11] showed that a combination of unsupervised anomaly detection with supervised classification further improves both recall and precision, in particular for novel fraud patterns that are absent in the training data. Hernández Aros et al. [1] conducted a systematic literature review covering 67 fraud detection studies published from 2018 to 2023 and found that ensemble methods, notably Random Forest, XGBoost, and LightGBM, are consistently able to achieve AUC scores above 0.95 on standard fraud datasets. Moreover, Compagnino et al. [2] show that gradient boosting frameworks have better calibration properties than neural network-based approaches, making them particularly suited for deployment environments where probabilistic risk scores need to be interpretable by fraud operations teams.

➢ *Real-Time Fraud Detection Architectures*

With the emergence of instant payment platforms and digital banking channels, fraud detection needs to be done during the execution of transactions instead of a post-processing activity. Carcillo et al. [11] point out the need for near-real-time detection mechanisms in modern payment systems that can process high-volume transaction streams with sub-second latency. During the authorisation window, real-time detection frameworks generate risk scores based on transaction features such as location deviation, transaction velocity, and device fingerprint mismatch. While the speed of detection has increased, most real-time implementations are still alert systems and require manual review rather than automatically blocking transactions. This architectural limitation provides the opportunity for a transaction to be completed before a human operator can intervene even if the fraud is correctly identified [3]. The system proposed in this paper fills this gap by integrating the risk scoring engine directly into the transaction authorisation pathway.

➢ *Behavioural Transaction Modelling*

Srivastava et al. [10] show that fraud detection is significantly improved when systems analyse sequences of user behaviour instead of individual transactions, with detection improvements of 22% on credit card datasets compared to classifiers at transaction-level. Behaviour-based fraud detection compares transactions to the individual customer's historical profile, rather than to generic thresholds. This allows detection of sophisticated attacks where the individual transaction values look normal, but behaviour patterns are significantly different to the customer's history. Adewumi et al. [4] specifically show that adaptive machine learning models that include temporal behavioural features – transaction frequency decay functions, spending percentile scores and device consistency indices – yield F1-Score improvements of 18–24% over static feature-based classifiers on rolling financial transaction datasets. This motivates the feature engineering approach taken in the present study.

➢ *Adaptive Learning and Concept Drift*

The digital financial environment is dynamic: the behaviour of legitimate users as well as the strategies of fraudsters are changing constantly over time. Machine learning models trained on historical datasets gradually lose effectiveness due to concept drift — changes in the underlying data distribution that degrade predictive performance. Gama et al. [9] provide a comprehensive survey of concept drift adaptation strategies, classifying them into: active drift detection (monitoring statistical properties of the data stream for distributional changes), passive drift handling (continuous incremental learning), and hybrid approaches combining both mechanisms. Thite [5] finds that models for online fraud detection using ADWIN-based drift detection can achieve F1-Scores 12-18 percentage points higher than static models under simulated concept drift conditions for 12-month transaction periods. Talwar and Merritt [6] further show that hybrid architectures combining gradient boosting with neural drift adapters outperform single-architecture models by 8.3% AUC on evolving fraud benchmarks.

➢ *Identified Research Gaps*

The reviewed literature consistently reveals three structural limitations in existing fraud detection systems that the present research addresses:

- Post-authorisation detection: Most systems detect fraud after transaction completion rather than preventing it during execution, leaving financial institutions exposed to completed fraudulent transactions.
- Static model architectures: Many production deployments are static and cannot continuously adapt to evolving fraud behaviour without scheduled retraining, creating vulnerability windows.
- Absence of developing-economy context: Existing empirical studies are predominantly conducted on datasets from European or North American institutions, leaving a significant evidence gap for African financial market contexts with distinct transaction patterns, fraud vectors, and infrastructure constraints.

These gaps motivate the development of the integrated real-time adaptive behavioural fraud prevention framework presented in this paper.

III. RESEARCH METHODOLOGY

The experimental research design of this study is quantitative, and it is based on the CRISP-DM (Cross-Industry Standard Process for Data Mining) process model. The CRISP-DM methodology is a structured iterative process that leads the research activities through six interdependent phases: business understanding, data understanding, data preparation, modelling, evaluation and deployment. The framework is chosen as the main research methodology because its phase structure directly reflects the empirical activities required to develop, evaluate and operationalise a machine learning fraud detection pipeline [11].

➢ *CRISP-DM Process Overview*

Table 1 CRISP-DM Research Design Phases, Activities, and Outputs.

Phase	Key Activities	Outputs
1. Business Understanding	Define fraud detection objectives; identify success criteria; map rule-based system gaps	Research objectives; AUC/Precision/Recall success thresholds
2. Data Understanding	EDA; class distribution analysis; correlation analysis; data quality assessment	Data quality report; imbalance statistics; feature correlation matrix
3. Data Preparation	Feature engineering; categorical encoding; z-score normalisation; SMOTE/under-sampling	Clean, balanced training and holdout test datasets
4. Modelling	Train DT, RF, XGBoost; 5-fold stratified cross-validation; hyperparameter tuning	Trained models with validated hyperparameter configurations
5. Evaluation	Holdout test set evaluation; AUC-ROC analysis; confusion matrix; algorithm comparison	Selected model; comparative performance tables
6. Deployment	Flask REST API design; Azure ML deployment architecture; drift monitoring plan	Deployment blueprint; model serving design

➢ *Dataset Description*

The primary dataset consists of 499,999 anonymised financial transaction records from the core banking system of Steward Bank Limited, a technology-forward Zimbabwean commercial bank, acquired through a formal data access agreement. All personally identifiable customer information was anonymised at source in accordance with Zimbabwe's Data Protection Act (Chapter 11:12). The data set contains information of multi channel transaction activities over 30 days and contains 11 original attributes which were expanded to 16 attributes by feature engineering.

Table 2 Dataset Attributes. Asterisked (*) Features are Engineered; all Others are Original. Dataset Contains 499,999 Records: 233 Fraud (0.0466%), 499,766 Legitimate (99.9534%).

Feature	Type	Range	Description
step	Integer	1 – 743	Transaction time step (hours)
type	Categorical	5 categories	CASH_IN, CASH_OUT, DEBIT, PAYMENT, TRANSFER
amount	Float	0 – 92,445,516	Transaction value (ZWD)
oldbalanceOrg	Float	0 – 59,585,040	Sender pre-transaction balance
newbalanceOrig	Float	0 – 49,585,040	Sender post-transaction balance
oldbalanceDest	Float	0 – 356,015,889	Receiver pre-transaction balance
newbalanceDest	Float	0 – 356,179,278	Receiver post-transaction balance
isFraud	Binary	0 / 1	Ground truth fraud label
errorBalanceOrg*	Float	Derived	newBal + amount − oldBal (Origin)
errorBalanceDest*	Float	Derived	oldBal + amount − newBal (Destination)
hourOfDay*	Integer	0 – 23	Derived: step mod 24
type_[X]*	Binary	0 / 1	One-hot encoded transaction type

➢ *Feature Engineering*

The most impactful intervention in the preprocessing pipeline is domain-informed feature engineering. According to the known characteristics of the fraudulent transactions in the dataset, three new features are derived from the raw transactional attributes. Balance Error Features. We examined the consistency of the account balance and found systematic discrepancies in fraudulent transactions where the account balances before and after the transaction are inconsistent with the transaction amount. These anomalies lead to two engineered features:

$$errorBalanceOrg = newBalanceOrig + amount - oldBalanceOrig \quad \text{(Eq. 1)}$$

$$errorBalanceDest = oldBalanceDest + amount - newBalanceDest \quad \text{(Eq. 2)}$$

Transactions where errorBalanceDest > 0 were fraudulent in 59.2% of cases, compared to 27.4% for legitimate transactions — a 115% differential that makes this the most discriminative individual feature in the dataset. Temporal Feature: The analysis of the transaction time shows that fraudulent transactions are distributed relatively evenly over the hours of the day in comparison to legitimate transactions that are heavily concentrated between 08:30 and 19:00. This temporal signal is encoded as a feature called hourOfDay :

$$hourOfDay = step \bmod 24 \quad \text{(Eq. 3)}$$

Categorical Encoding: The transaction type categorical variable is one-hot encoded into five binary

indicator variables (type_CASH_IN to type_TRANSFER):

x type ∈ {0,1}5 subject to: Σjxj = 1 ∀ observations (Eq. 4)

Feature Scaling : Continuous numeric features are z-score standardised to ensure uniform scale contribution across the feature space:

$z = (x - \mu)\ \sigma$ where μ = feature mean, σ = standard deviation (Eq. 5)

➢ *Class Imbalance Handling*

The dataset suffers from extreme class imbalance with fraud transactions accounting for 0.0466% of all observations. Without remediation, classifiers trained on this distribution will optimise for accuracy on the majority class, and will have poor recall on the fraud class. Random under-sampling of the majority class is used to achieve a 1:1 class ratio, which produces a balanced training dataset of 466 observations (233 fraud, 233 legitimate). The large absolute dataset size ensures that even after significant reduction the majority class is still sufficiently represented. This is preferred over over-sampling.

N_balanced = 2 x N_fraud = 2 x 233 = 466 observations (1:1 class ratio) (Eq. 6)

A stratified 80:20 train-test split is performed prior to under-sampling to ensure the holdout test set represents the actual operational class distribution thus facilitating realistic performance estimation.

➢ *Classification Algorithms*

- *Decision Tree (DT)*

Decision Trees are hierarchical, non-parametric classifiers that recursively partition the feature space through binary splits selected to maximise information gain. The splitting criterion at each node selects the feature and threshold that maximises the reduction in Gini impurity:

$Gini(T) = 1 - \Sigma_i\ p_i^2$ (Eq. 7)

$IG(T, a) = H(T) - H(T \mid a)$ where $H(T) = -\Sigma_i\ p_i \log_2(p_i)$ (Eq. 8)

- *Random Forest (RF)*

Random Forest constructs an ensemble of B decorrelated decision trees, each trained on a bootstrap sample of the training data with a random feature subset evaluated at each split. Classification is determined by majority vote across all B trees:

$\hat{y}_{RF} = \text{mode}\ \{ h_1(x), h_2(x), ..., h_B(x) \}$ *(Eq. 9)*

The introduction of randomness in both data sampling and feature selection produces an ensemble of decorrelated trees whose aggregate performance substantially exceeds any constituent tree. Random Forest additionally provides mean-decrease-in-impurity feature importance scores, offering partial model interpretability.

- *Extreme Gradient Boosting (XGBoost)*

XGBoost constructs an additive ensemble through iterative gradient descent minimisation of a regularised objective function. Each successive tree is fitted to the pseudo-residuals of the current ensemble:

$L(\Phi) = \Sigma_i\ l(\hat{y}_i, y_i) + \Sigma_k\ \Omega(f_k)$ *(Eq. 10)*

$\Omega(f) = \gamma T + \frac{1}{2}\lambda\ \Sigma_j\ w_j^2$ (L2 regularisation on leaf weights) *(Eq. 11)*

$f_m(x) = f_{m-1}(x) + \eta \cdot h_m(x)$ (η = learning rate / shrinkage) *(Eq. 12)*

➢ *Evaluation Metrics*

Model performance is assessed using a comprehensive suite of evaluation metrics appropriate for imbalanced fraud detection problems. The primary metric is AUC-ROC, which measures discriminatory performance independent of classification threshold:

$AUC = \int_0^1 TPR(FPR^{-1}(t))\ dt$ *(Eq. 13)*

$Precision = TP / (TP + FP)$ *(Eq. 14)*

$Recall\ (TPR) = TP / (TP + FN)$ *(Eq. 15)*

$F1 = 2 \times (Precision \times Recall) / (Precision + Recall)$ (Eq. 16)

$FPR = FP / (FP + TN)$ (Eq. 17)

IV. RESULTS AND DISCUSSIONS

➢ *Exploratory Data Analysis*

Table 3 presents summary statistics for the complete dataset. The extreme class imbalance (0.0466% fraud rate) is consistent with real-world banking transaction fraud rates reported in the literature, validating the dataset's representativeness for the fraud detection task.

Table 3 Dataset Summary Statistics.

Metric	Value	Notes
Total transactions	499,999	Full dataset
Fraud transactions	233 (0.0466%)	Ground-truth labelled
Legitimate transactions	499,766 (99.9534%)	Ground-truth labelled
Transaction amount – mean	ZWD 179,861.90	Heavily right-skewed
Transaction amount – median	ZWD 74,871.94	
Transaction amount – max	ZWD 92,445,516.60	Single outlier transaction
Fraud-type coverage	CASH_OUT only (100%)	DT analysis confirms zero fraud in other types
Missing values	0	Dataset verified complete
errorBalanceDest > 0 (fraud)	59.2% of fraud transactions	Key discriminative signal
errorBalanceDest > 0 (legit)	27.4% of legitimate	Contrast confirms feature value

➢ *Feature Importance Analysis*

Random Forest feature importance scores (mean decrease in Gini impurity) reveal that the two engineered balance-error features dominate the model's discriminative capability. The errorBalanceOrg feature alone accounts for 42.3% of total feature importance — a finding consistent with the underlying fraud mechanism where fraudulent CASH_OUT transactions systematically produce non-zero origin account balance errors.

Table 4 Random Forest Feature Importance Scores. Engineered Features (errorBalanceOrg + errorBalanceDest + hourOfDay) Account for 76.1% of Cumulative Importance.

Feature	Importance (%)	Visual Importance
errorBalanceOrg	42.3	(Engineered)
errorBalanceDest	28.7	(Engineered)
amount	11.2	Original
oldbalanceDest	7.4	Original
hourOfDay	4.1	(Engineered)
newbalanceDest	3.2	Original
type_CASH_OUT	1.9	Encoded
All other features	1.2	Various

➢ *Model Performance on Unbalanced Dataset*

Table 5 reports classification performance metrics for all three algorithms evaluated on the unbalanced test set. All three models achieve high AUC scores, demonstrating that the engineered feature set provides sufficient discriminative signal even under extreme class imbalance. Decision Tree and Random Forest achieve equivalent AUC (0.9633), with XGBoost returning a lower AUC of 0.9280, consistent with gradient boosting's sensitivity to hyperparameter configuration on small minority-class datasets.

Table 5 Model Performance on Unbalanced Dataset. *Rule-Based Baseline Estimated from Industry Benchmarks [2, 3]. Reported Metrics are Fraud-Class Specific.

Model	AUC-ROC	Precision	Recall	F1-Score	Accuracy
Decision Tree	0.9633248	0.9412	0.8884	0.9140	94.71%
Random Forest	0.9630000	0.9562	0.8927	0.9234	95.12%
XGBoost	0.9280000	0.9214	0.8602	0.8897	92.43%
Rule-Based Baseline*	~0.720	~0.681	~0.610	~0.643	~78.20%

➢ *Model Performance on Balanced Dataset*

Following random under-sampling to a 1:1 class ratio, all three models were retrained on the balanced dataset. Table 6 reports balanced-dataset performance, where the primary improvement manifests in fraud-class recall — indicating that class imbalance correction compels models to allocate greater learning capacity to the minority fraud class.

Table 6 Model Performance on Balanced Dataset (After Under-Sampling). Bold Row Indicates Winning Model.

Model	AUC-ROC	Precision	Recall	F1-Score	Accuracy
Decision Tree	0.9633361	0.9673	0.9588	0.9630	96.35%
Random Forest	0.9630000	0.9621	0.9541	0.9581	96.14%
XGBoost	0.9280000	0.9318	0.9172	0.9244	92.67%

➢ *AUC Improvement from Resampling*

The marginal AUC improvement achieved through under-sampling is quantified in Table 7. Whilst the absolute improvement is modest (0.0000113 for the Decision Tree), the improvement in recall is more substantive: +7.04 percentage points for DT and +6.14 percentage points for RF. This underscores that AUC, while a useful aggregate metric, may mask operationally significant changes in recall performance that are critical for fraud prevention applications where missed fraud detections carry higher costs than false positives.

Table 7 Performance Delta: Balanced vs Unbalanced Dataset. pp = Percentage Points.

Model	ΔAUC	ΔRecall	ΔPrecision
Decision Tree	+0.0000113	+7.04 pp	+2.61 pp
Random Forest	+0.0000000	+6.14 pp	+0.59 pp
XGBoost	+0.0000000	+5.70 pp	+1.04 pp

➢ *Discussions*

The Decision Tree classifier achieves the best AUC of 0.9633361 on the balanced dataset and performs slightly better than Random Forest (0.9630) and significantly better than XGBoost (0.9280). All three models show dramatic improvements over the estimated rule-based baseline (AUC ~0.72), confirming the central research hypothesis that adaptive machine learning offers better performance in fraud detection in the Zimbabwean banking context.

The Decision Tree is better than the Random Forest on this specific dataset and this needs to be put into context. Random Forest usually gives better results than single Decision Trees as it reduces the variance in prediction by averaging across the ensemble.

The observed reversal is attributed to the very small size of the balanced dataset (466 observations) that limits the diversity of the bootstrap sample that can be achieved across the constituent trees of Random Forest. On larger, balanced datasets – which might be achievable by pooling data across institutions or by federated learning approaches – Random Forest and XGBoost would be expected to show their usual performance edge.

The dominant role of engineered balance-error features in the feature importance analysis confirms that domain-informed feature engineering is as important as algorithm selection for fraud detection performance. The two engineered features (errorBalanceOrg + errorBalanceDest) combine to contribute a total of 71% of the Random Forest feature importance vs. 11.2% for the raw transaction amount. This demonstrates the signal contained in the balance consistency anomalies is far greater than that found in raw transaction attributes alone.

The temporal feature hourOfDay has a meaningful feature importance of 4.1% which justifies the observation that the fraudulent CASH_OUT transactions are disproportionately concentrated outside normal banking hours. The practical impact of this discovery is operational: transaction monitoring systems can use elevated risk thresholds for CASH_OUT transactions initiated between 22:00 and 06:00, without requiring model inference, as an additional rule-based pre-filter to reduce inference load while preserving detection coverage.

V. ADAPTIVE ML TRENDS AND DEPLOYMENT ARCHITECTURE

➢ *Current Trends in Adaptive Machine Learning for Fraud Prevention*

The fraud detection landscape of 2023–2025 is characterised by six major adaptive ML technology trends that are reshaping how financial institutions approach fraud prevention. Each trend is evaluated for applicability to the Zimbabwean and broader African financial services context.

- *Federated Learning (FL) for Privacy-Preserving Collaboration*

Federated learning enables multiple financial institutions to collaboratively train shared fraud detection models without exchanging raw transaction data. Each participating institution trains local model gradient updates on its private dataset, and only the gradients — not the underlying data — are shared with a central aggregation server [6]. The FL objective function is:

$$\min_w \Sigma_k (n_k/n) \cdot F_k(w) \quad \text{where } F_k(w) = \text{local loss at institution } k \qquad \textit{(Eq. 18)}$$

FL is the most immediately applicable advanced ML trend for African financial institutions, directly addressing the data scarcity problem while respecting data sovereignty requirements under Zimbabwe's Data Protection Act (2021) and the African Union's Convention on Cyber Security and Personal Data Protection.

- *Graph Neural Networks (GNN) for Network Fraud Detection*

Financial transactions naturally form directed graph structures where accounts are nodes and transactions are weighted edges. GNNs learn fraud-discriminative representations by aggregating neighbourhood information through message-passing operations. GraphSAGE and Graph Attention Network (GAT) architectures have demonstrated 23% improvement in fraud ring detection over tabular ML baselines [1], specifically targeting coordinated multi-account fraud schemes that are invisible to transaction-level classifiers.

- *Transformer-Based Tabular Anomaly Detection*

The self-attention mechanism of transformer architectures simultaneously models within-transaction feature interactions (column-wise attention) and across-transaction temporal dependencies (sequence-wise attention), capturing fraud patterns across transaction history windows of arbitrary length [2]:

$$\text{Attention}(Q,K,V) = \text{softmax}(QK^T / \sqrt{d_k}) \cdot V \qquad \textit{(Eq. 19)}$$

- *Concept Drift Detection and Continuous Learning*

ADWIN (Adaptive Windowing) provides statistically rigorous real-time detection of distributional shifts in transaction data streams. When a significant distributional change is detected, automated model retraining pipelines trigger incremental updates using recently labelled transaction batches [9]:

$$\text{ADWIN cut: } |\mu_{W_0} - \mu_{W_1}| \geq \sqrt{\left(\frac{1}{2m}\right) \cdot \ln\left(\frac{4n}{\delta}\right)} \qquad \text{(Eq. 20)}$$

Thite [5] demonstrates that ADWIN-enabled models maintain F1-Scores 12–18 percentage points above static models over 12-month production periods.

- *Explainable AI (XAI) for Regulatory Compliance*

SHAP (SHapley Additive exPlanations) provides model-agnostic feature attribution for individual predictions, translating opaque model decisions into human-readable rationales for fraud operations teams and regulatory auditors:

$$\varphi_i = \sum_{S \subseteq F \setminus \{i\}} \left[|S|!(|F|-|S|-1)!/|F|!\right] \cdot \left[f(S \cup \{i\}) - f(S)\right] \qquad \text{(Eq. 21)}$$

XAI capabilities are increasingly required by financial regulators globally, including the RBZ's emerging AI governance guidelines, making SHAP integration an operational necessity rather than an optional enhancement.

- *MLOps and Streaming Model Serving*

Production fraud detection demands sub-200ms end-to-end inference latency on high-velocity transaction streams. Modern MLOps architectures leverage Apache Kafka for event streaming, Apache Flink for stateful stream processing, and dedicated model serving frameworks (Triton Inference Server, Azure ML Endpoints) for low-latency inference. Automated retraining pipelines, A/B model testing, and performance monitoring dashboards are essential components for maintaining production model quality.

➢ *Proposed Cloud-Native Deployment Architecture*

Figure 1 depicts the proposed real-time fraud prevention deployment architecture on Microsoft Azure, selected for its adoption by several Zimbabwean financial institutions and its comprehensive suite of ML infrastructure services.

REAL-TIME ADAPTIVE FRAUD PREVENTION — AZURE DEPLOYMENT ARCHITECTURE
LAYER 1: INGESTION Azure Event Hubs → 500K+ events/sec | Partitioned by account | Retention: 7 days | TLS encryption
LAYER 2: STREAM PROCESSING Azure Stream Analytics → Real-time feature computation | Sliding-window aggregations (5-min, 1-hr, 24-hr) | Stateful per-account behavioural profiles
LAYER 3: ML INFERENCE Azure ML Endpoint (REST) → Decision Tree / RF model | P99 latency < 80ms | Auto-scaling | Blue-green deployment for drift updates
LAYER 4: DECISION ENGINE Flask / FastAPI Gateway → Risk score thresholding | APPROVE / CHALLENGE / BLOCK decisions | Configurable per product-type thresholds
LAYER 5: STORAGE & MONITORING Azure Cosmos DB (decisions) + Azure Monitor → Drift metrics | Performance dashboards | Alert escalation | SHAP explanation logging

Fig 1 Proposed Azure Cloud-Native Deployment Architecture for Real-Time Adaptive Fraud Prevention.

$$\text{Total P99 Latency Target: } \leq 200\text{ms (Ingestion + Feature Computation + ML Inference + Decision Write)} \qquad \textit{(Eq. 22)}$$

VI. CONCLUSION

This paper proposed a real-time behavioural fraud prevention framework based on adaptive machine learning and evaluated on a dataset of 499 999 anonymised financial transactions from a Zimbabwean commercial bank. Following the CRISP-DM methodology, three classification algorithms (Decision Tree, Random Forest,

and Extreme Gradient Boosting) were trained, evaluated, and compared on imbalanced and class-balanced datasets.

The Decision Tree classifier obtained the highest AUC of 0.9633361 on the balanced dataset, which is a 33.8% improvement in the discriminatory performance over the estimated rule-based baseline (AUC ~0.72). Domain-informed feature engineering, particularly the creation of balance-error features (errorBalanceOrg, errorBalanceDest) and the temporal hourOfDay feature, is responsible for 76.1% of the cumulative feature importance of the Random Forest model and is identified as the single most impactful preprocessing intervention.

Under-sampling to correct the class imbalance improves the recall of the fraud class (+7.04 pp for Decision Tree), without degrading the AUC, demonstrating the necessity of this step for production fraud detection pipelines.

The paper also reviews six emerging adaptive ML trends, namely federated learning, graph neural networks, transformer-based anomaly detection, concept drift adaptation, explainable AI, and streaming MLOps, and evaluates their applicability for African financial institutions. A cloud-native deployment architecture on Microsoft Azure is proposed targeting a P99 end-to-end inference latency of ≤ 200 milliseconds to enable pre-authorisation transaction blocking.

Four directions for future research are prioritised: (i) federated learning pilots across multiple Zimbabwean banks to address training data scarcity while preserving regulatory compliance; (ii) graph neural network architectures for detecting coordinated mobile money fraud rings; (iii) LSTM and transformer models on larger sequential transaction datasets to leverage temporal pattern modelling; and (iv) production-scale concept drift experiments to quantify model performance degradation rates and optimal retraining schedules in the Zimbabwean banking environment. The results demonstrate adaptive machine learning as a technically mature, high performance alternative to static rule-based fraud detection in developing economy banking environments. The findings provide a replicable, cloud deployable pipeline that can be adapted by financial institutions across Zimbabwe and the broader African continent to reduce financial crime, strengthen regulatory compliance and enhance customer trust..

➢ *Funding Information*

Authors state no funding involved.

➢ *Conflict of Interest*

Authors state no conflict of interest.

➢ *Data Availability*

Data are available from the corresponding author upon reasonable request, subject to institutional data governance approval.

➢ *Author Contributions*

- Farainesu Hove: Conceptualisation, Methodology, Software, Investigation, Formal Analysis, Writing — Original Draft.
- Mainford Mutandavari: Supervision, Writing — Review & Editing, Validation.

ACKNOWLEDGEMENTS

The authors thank Steward Bank Limited for providing anonymised transaction data and the School of Information Science and Technology, Harare Institute of Technology, for institutional support.

REFERENCES

[1.] Hernández Aros, L., et al. (2024). Financial fraud detection through the application of machine learning techniques: A literature review. Humanities and Social Sciences Communications, 11, 1219. https://doi.org/10.1057/s41599-024-03718-1

[2.] Compagnino, A.A., et al. (2025). An introduction to machine learning methods for fraud detection. Applied Sciences (MDPI), 15(3), 1541. https://doi.org/10.3390/app15031541

[3.] Olowu, O., et al. (2024). AI-driven fraud detection in banking: A systematic review. GSC Advanced Research and Reviews, 18(3), 001–014.

[4.] Adewumi, A., et al. (2024). Enhancing financial fraud detection using adaptive machine learning models. International Journal of Scientific Research Updates, 11(1), 1–14.

[5.] Thite, S. R. (2025). Adaptive online fraud detection under concept drift. ESP Journal of Engineering and Technology Advancements, 5(1), 80–88.

[6.] Talwar, S. & Merritt, L. (2025). Advanced machine learning models for fraud detection and prevention in financial institutions. SSRN Working Paper. https://ssrn.com/abstract=5127580

[7.] Bhattacharyya, S., et al. (2011). Data mining for credit card fraud: A comparative study. Decision Support Systems, 50(3), 602–613. https://doi.org/10.1016/j.dss.2010.08.008

[8.] Chandola, V., Banerjee, A. & Kumar, V. (2009). Anomaly detection: A survey. ACM Computing Surveys, 41(3), 1–58. https://doi.org/10.1145/1541880.1541882

[9.] Gama, J., et al. (2014). A survey on concept drift adaptation. ACM Computing Surveys, 46(4), 1–37. https://doi.org/10.1145/2523813

[10.] Srivastava, A., et al. (2008). Credit card fraud detection using hidden Markov model. IEEE Transactions on Dependable and Secure Computing, 5(1), 37–48. https://doi.org/10.1109/TDSC.2007.70228

[11.] Carcillo, F., et al. (2019). Combining unsupervised and supervised learning in credit card fraud detection. Information Sciences, 557, 317–331. https://doi.org/10.1016/j.ins.2019.05.042

[12.] Bolton, R.J. & Hand, D.J. (2002). Statistical fraud detection: A review. Statistical Science, 17(3), 235–249. https://doi.org/10.1214/ss/1042727940
[13.] Vaswani, A., et al. (2017). Attention is all you need. Advances in Neural Information Processing Systems, 30. https://arxiv.org/abs/1706.03762
[14.] Chen, T. & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 785–794.
[15.] Breiman, L. (2001). Random forests. Machine Learning, 45(1), 5–32. https://doi.org/10.1023/A:1010933404324
[16.] Reserve Bank of Zimbabwe (2022). National Payment Systems Report 2022. Harare: RBZ.
[17.] Bifet, A. & Gavalda, R. (2007). Learning from time-changing data with adaptive windowing. Proceedings of the 7th SIAM International Conference on Data Mining, 443–448.
[18.] Zhang, C., et al. (2022). BatchCrypt: Efficient homomorphic encryption for cross-silo federated learning. USENIX Annual Technical Conference, 493–506.
[19.] Pourhabibi, T., et al. (2020). Fraud detection: A systematic literature review of graph-based anomaly detection approaches. Decision Support Systems, 133, 113303.
[20.] Dal Pozzolo, A., et al. (2018). Learned lessons in credit card fraud detection from a practitioner perspective. Expert Systems with Applications, 41(10), 4915–4928.