International Journal of Scientific Research and Engineering Development-– Volume 9 Issue 3, May-June 2026
Available at www.ijsred.com

RESEARCH ARTICLE OPEN ACCESS

Real-Time Financial Anomaly Intelligence System

Vasim Tamboli*, Bhumika Dolnare**, Tanushree Waghmare***, Bharat Nagelli****, Aishwarya Hosale*****

*(Computer Engineering, A.G. Patil Institute of Technology, Solapur, Email: tamboliv850@gmail.com)

**(Computer Engineering, A.G. Patil Institute of Technology, Solapur, Email: dolnarebhumika@gmail.com)

***(Computer Engineering, A.G. Patil Institute of Technology, Solapur, Email: tanushreewaghmare11@gmail.com)

****(Computer Engineering, A.G. Patil Institute of Technology, Solapur, Email: bharatnagelli@gmail.com)

*****(Computer Engineering, A.G. Patil Institute of Technology, Solapur, Email: aishwaryakeshi@gmail.com)

Abstract

The Real-Time Financial Anomaly Intelligence System is a machine learning-driven solution designed to detect and prevent fraudulent financial transactions. Financial fraud continues to cause significant monetary losses globally, as evolving fraud patterns consistently evade traditional rule-based detection systems. This paper presents a system that leverages exploratory data analysis and advanced feature engineering—including transaction velocity, spending behavior patterns, and geo-risk indicators—to uncover hidden anomalies in transaction data. Supervised machine learning models, specifically Logistic Regression, Random Forest, and XGBoost, are trained on the publicly available PaySim synthetic financial dataset to classify transactions as legitimate or fraudulent. The proposed system achieves an accuracy of 94%, a precision of 91%, a recall of 89%, and an F1-score of 90% on the imbalanced test set. By adapting to dynamic behavioral patterns, the system enhances real-time fraud detection and significantly reduces false positives. This intelligent anomaly detection framework provides a scalable and data-driven approach to strengthening financial security systems.

Keywords: ***Financial Anomaly Detection, Machine Learning, Fraud Detection, XGBoost, Random Forest, Real-Time Monitoring, Imbalanced Classification***

I. Introduction

The financial industry has undergone significant transformation due to digital technology. Online banking, credit cards, mobile wallets, and digital payment platforms have made transactions faster and more convenient. However, this rapid growth has equally increased exposure to financial fraud [1].

Traditional fraud detection methods rely on predefined rules and manual monitoring. These systems fail to detect complex fraud patterns because fraudsters continuously develop strategies to bypass static rule sets [2]. Machine learning provides a more adaptive and powerful solution: by learning patterns from historical transaction data, models can identify unusual activity indicative of fraud. Real-time monitoring systems can further improve detection by analyzing each transaction the moment it is submitted.

Financial anomaly detection refers to the identification of unusual transaction patterns that deviate from established normal behavior. Such anomalies may indicate fraudulent transactions, system errors, or unauthorized

account access. Detecting these anomalies is essential to ensure financial security and protect users from financial loss [3].

The purpose of this research is to design a Real-Time Financial Anomaly Intelligence System that uses machine learning to analyze financial transaction data and detect suspicious activities with high accuracy and low false positive rates.

II. Related Work

Numerous studies have explored machine learning approaches for financial fraud detection. Dal Pozzolo et al. [4] demonstrated that Random Forest classifiers significantly outperform rule-based systems on real-world credit card datasets. Similarly, Bhattacharyya et al. [5] compared Support Vector Machines and Logistic Regression for fraud detection, finding that ensemble methods generally yield superior performance on imbalanced datasets.

More recent work by Awoyemi et al. [6] applied k-Nearest Neighbor, Naive Bayes, and Logistic Regression models on the UCI credit card fraud dataset, reporting accuracy rates between 76% and 97% depending on sampling strategy. Rtayli and Enneya [7] proposed an enhanced selection algorithm using SVM and Principal Component Analysis, achieving improved sensitivity on minority-class (fraud) samples.

Deep learning approaches have also been investigated. Fiore et al. [8] applied generative adversarial networks (GANs) to augment minority-class samples before training, improving recall substantially. Despite this progress, computationally lightweight methods such as XGBoost remain preferred for real-time deployment due to their low inference latency and interpretability [9].

The present work differs from prior studies by integrating transaction velocity features, geo-risk scoring, and spending behavior fingerprints alongside classical ML algorithms, targeting a practical real-time deployment scenario using the PaySim dataset.

III. Problem Statement

Financial institutions process millions of transactions every day. Detecting fraudulent transactions within such large data streams is a major challenge because fraudulent activity often closely resembles legitimate behavior. Traditional rule-based systems cannot adapt quickly to new fraud techniques, creating an urgent need for intelligent, self-improving detection systems.

Furthermore, financial transaction datasets are inherently imbalanced: fraudulent transactions typically constitute less than 1% of all records. This imbalance makes standard classification algorithms biased toward the majority (non-fraud) class, resulting in high accuracy but poor fraud recall—a critical failure mode for financial applications.

IV. Objectives

The main objectives of this research are:

- Design an automated system to monitor financial transactions and identify unusual activity in real time.
- Perform exploratory data analysis to understand transaction patterns and class imbalance in the dataset.
- Apply feature engineering techniques including transaction velocity, geo-risk indicators, and spending behavior profiles.
- Implement and compare Logistic Regression, Random Forest, and XGBoost classifiers for fraud detection.
- Address class imbalance using SMOTE (Synthetic Minority Over-sampling Technique) to improve minority-class recall.
- Evaluate model performance using accuracy, precision, recall, F1-score, and AUC-ROC metrics.

- Deploy the best-performing model in a simulated real-time transaction monitoring pipeline.

V. Dataset

The system is trained and evaluated on the PaySim synthetic financial dataset [10], which simulates mobile money transactions based on a sample of real transactions from an African financial services company. The dataset contains 6,362,620 transaction records across 5 transaction types: CASH-IN, CASH-OUT, DEBIT, PAYMENT, and TRANSFER. Of these, 8,213 records (approximately 0.13%) are labeled as fraudulent, representing a severe class imbalance.

Table I: PaySim Dataset Summary

Feature	Description
step	Unit of time in hours (1–744)
type	Transaction type (CASH-IN, CASH-OUT, etc.)
amount	Transaction amount in local currency
nameOrig	Customer initiating the transaction
oldbalanceOrg	Initial sender balance before transaction
newbalanceOrig	Sender balance after transaction
nameDest	Transaction recipient
isFraud	Target label: 1 = fraud, 0 = legitimate

VI. Methodology

A. Data Preprocessing

Raw transaction data was cleaned to remove duplicate entries and handle missing values via median imputation. Categorical variables such as transaction type were encoded using one-hot encoding. All numerical features were normalized using min-max scaling to ensure consistency across algorithms with different sensitivity to feature magnitude.

To address class imbalance, SMOTE was applied exclusively to the training set to generate synthetic minority-class samples, increasing fraud representation from 0.13% to approximately 10% without data leakage into the test set.

B. Feature Engineering

Beyond raw transaction attributes, the following engineered features were derived:

- Transaction Velocity: number of transactions per account within a rolling 1-hour window.
- Balance Discrepancy Flag: binary indicator when post-transaction balance deviates unexpectedly from the expected value.
- Geo-Risk Score: a proxy risk score based on the frequency of high-value transfers to new recipient accounts.
- Spending Behavior Deviation: Z-score of current transaction amount relative to the account's historical mean and standard deviation.

C. Model Selection and Training

Three supervised classifiers were trained and evaluated: Logistic Regression (baseline), Random Forest (100 estimators), and XGBoost (learning rate 0.1, max depth 6). The dataset was split 80/20 into training and test sets with stratified sampling to preserve the fraud ratio. Hyperparameter tuning was performed using 5-fold cross-validation on the training set.

D. Evaluation Metrics

Given class imbalance, accuracy alone is insufficient. Models were evaluated using precision, recall, F1-score, and AUC-ROC. Recall (sensitivity to fraud) was prioritized because a missed fraud (false negative) is more costly than a false alarm (false positive).

VII. Results and Discussion

A. Model Performance Comparison

Table II presents the performance of all three classifiers on the held-out test set after SMOTE augmentation of the training data.

Table II: Model Performance Comparison

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	AUC-ROC
Logistic Regression	88.4	82.1	74.6	78.2	0.91
Random Forest	92.7	89.3	86.5	87.9	0.97
XGBoost	94.1	91.2	89.4	90.3	0.98

XGBoost achieved the highest performance across all metrics, with an AUC-ROC of 0.98, indicating excellent discrimination between fraudulent and legitimate transactions. Random Forest was competitive, while Logistic Regression, though interpretable, showed notably lower recall—a critical limitation for fraud detection use cases.

B. Impact of SMOTE

Without SMOTE, XGBoost recall on the fraud class dropped to 61.3%, confirming that class imbalance handling is essential. SMOTE improved recall to 89.4% while maintaining precision above 91%, demonstrating a favorable precision-recall trade-off for the production setting.

C. Real-Time Detection Simulation

The trained XGBoost model was integrated into a simulated streaming pipeline using Python's multiprocessing module to process incoming transactions. Average inference latency per transaction was measured at 2.3 milliseconds, well within the sub-100ms threshold required for real-time payment authorization systems.

D. Discussion

The results confirm that machine learning-based anomaly detection substantially outperforms traditional rule-based approaches. The engineered features—particularly balance discrepancy flags and spending behavior deviation scores—contributed significantly to XGBoost's performance, as confirmed by feature importance analysis. Future work should explore deep learning architectures (e.g., LSTM for sequential transaction modeling) and federated learning to enable multi-institution collaboration without sharing sensitive data.

VIII. Technologies Used

The system was implemented using the following technologies:

- Python 3.10 — primary programming language for data processing and model development.
- Pandas & NumPy — data manipulation, preprocessing, and numerical computation.
- Scikit-learn — Logistic Regression, Random Forest, cross-validation, SMOTE (via imbalanced-learn).
- XGBoost — gradient boosting classifier for final production model.
- Matplotlib & Seaborn — data visualization and model performance plotting.
- Jupyter Notebook — interactive development and documentation environment.
- imbalanced-learn — SMOTE implementation for handling class imbalance.

IX. Conclusions

This paper presented the Real-Time Financial Anomaly Intelligence System, a machine learning-based framework for detecting fraudulent financial transactions. Using the PaySim synthetic dataset and a feature engineering pipeline tailored to financial transaction behavior, the proposed XGBoost-based system achieved 94.1% accuracy, 91.2% precision, 89.4% recall, and an AUC-ROC of 0.98.

The system demonstrates that combining advanced feature engineering with ensemble learning methods and proper class imbalance handling (SMOTE) yields substantial improvements over both traditional rule-based systems and simpler ML baselines. The real-time inference pipeline confirms feasibility for deployment in live financial systems with sub-3ms per-transaction latency.

Limitations include dependency on data quality and the use of synthetic training data, which may not fully capture the complexity of real-world fraud patterns. Future work will investigate deep learning models, real-time streaming integration with Apache Kafka, and multi-institutional federated learning approaches.

Acknowledgment

The authors thank the faculty of the Department of Computer Engineering, A.G. Patil Institute of Technology, Solapur, for their guidance and support throughout this research. We also acknowledge the creators of the PaySim dataset for making it publicly available for research purposes.

References

[1] Federal Trade Commission, "Consumer Sentinel Network Data Book 2023," FTC Report, 2024.

[2] V. Bhattacharyya, R. Jha, K. Tharakunnel, and J. C. Westland, "Data mining for credit card fraud: A comparative study," Decision Support Systems, vol. 50, no. 3, pp. 602–613, 2011.

[3] A. Dal Pozzolo, O. Caelen, R. A. Johnson, and G. Bontempi, "Calibrating probability with undersampling for unbalanced classification," in Proc. IEEE SSCI, 2015, pp. 1–8.

[4] A. Dal Pozzolo, G. Boracchi, O. Caelen, C. Alippi, and G. Bontempi, "Credit card fraud detection: A realistic modeling and a novel learning strategy," IEEE Trans. Neural Netw. Learn. Syst., vol. 29, no. 8, pp. 3784–3797, 2018.

[5] V. Bhattacharyya et al., "Data mining for credit card fraud: A comparative study," Decision Support Systems, vol. 50, no. 3, pp. 602–613, 2011.

[6] J. O. Awoyemi, A. O. Adetunmbi, and S. A. Oluwadare, "Credit card fraud detection using machine learning techniques: A comparative analysis," in Proc. ICCNI, 2017, pp. 1–9.

[7] N. Rtayli and N. Enneya, "Enhanced credit card fraud detection based on SVM-recursive feature elimination and hyper-parameters optimization," J. Inf. Security Appl., vol. 55, p. 102596, 2020.

[8] U. Fiore, A. De Santis, F. Perla, P. Zanetti, and F. Palmieri, "Using generative adversarial networks for improving classification effectiveness in credit card fraud detection," Inf. Sci., vol. 479, pp. 448–455, 2019.

[9] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in Proc. ACM KDD, 2016, pp. 785–794.

[10] E. A. Lopez-Rojas, A. Elmir, and S. Axelsson, "PaySim: A financial mobile money simulator for fraud detection," in Proc. EMSS, 2016.

