



A MACHINE LEARNING FRAMEWORK FOR REAL-TIME ANOMALY DETECTION IN FINANCIAL TRANSACTIONS

R.Rashika

Scholar,

*Department of Computer Application,
A.V.V.M. Sri Pushpam College (Autonomous),
Tiruvallur, Tamil Nadu, India.*

Dr. D. Ragupathi

Head of the Department,

*Department of Computer Applications,
A.V.V.M. Sri Pushpam College (Autonomous),
Tiruvallur, Tamil Nadu, India.*

Abstract

The rapid growth of digital transactions in sectors such as banking, e-commerce, and financial services has increased the risk of fraudulent activities and abnormal transaction behavior. Traditional fraud detection systems often rely on static, rule-based methods that fail to capture the evolving and sophisticated nature of modern anomalies. This paper proposes an intelligent framework for Real-Time Anomaly Detection in Transaction Data leveraging advanced machine learning algorithms. By analyzing continuous streams of transaction data, the system identifies suspicious patterns and outliers with high precision and minimal latency. The framework integrates feature engineering, adaptive learning, and real-time visualization to provide actionable insights for fraud prevention and risk management. Experimental evaluations conducted on simulated transaction streams demonstrate the system's effectiveness in maintaining high detection accuracy while adapting to dynamic data shifts. This research provides a scalable and robust solution for securing digital

financial ecosystems in high-velocity environments.

Keywords: Anomaly Detection, Machine Learning, Transaction Data, Real-Time Analytics, Fraud Detection, Financial Security, Streaming Data.

1. Introduction

The global digital economy has undergone a massive transformation, with electronic payments becoming the primary method for commerce. Distributed networks now support critical applications ranging from cloud computing and the Internet of Things (IoT) to complex enterprise financial systems. As these systems grow in scale and complexity, ensuring secure and reliable access control and transaction integrity becomes a paramount challenge for organizations. Digital platforms generate massive volumes of interaction data that contain vital indicators of system health and security.

Traditional anomaly detection mechanisms often suffer from significant limitations, such as single points of failure,



scalability issues, and vulnerability to sophisticated cyberattacks. These conventional systems are frequently reactive, analyzing data in batch mode days or weeks after a transaction has occurred. This delay allows fraudulent activities to persist, leading to massive revenue loss and escalating long-term recovery costs for both businesses and consumers. There is a critical need for systems that can monitor and analyze transaction activities in real time.

The primary motivation for this research is to leverage machine learning and streaming data technologies to overcome the limitations of traditional batch-based systems. Modern digital environments require decentralized and tamper-resistant mechanisms to ensure secure resource sharing and transaction validation. By moving away from periodic analysis, the proposed system ingests live data points—such as transaction records and access logs—and processes them through an incremental learning pipeline. This allows the models to evolve alongside shifting behavioral patterns, ensuring detection remains accurate even as market conditions change.

This project focuses on the design and implementation of a Blockchain-Based Secure Access Control system integrated with a Real-Time Anomaly Detection framework. The system utilizes a decentralized ledger to ensure transparency and immutability, while smart contracts automate the enforcement of security policies. In parallel, the anomaly detection engine continuously evaluates transaction features to flag suspicious

activities as they emerge. This dual-layered approach enhances resistance to attacks and prevents unauthorized access to sensitive financial resources.

Ultimately, the aim of this research is to provide a proactive and data-driven engagement model for financial security. By detecting early warning signs of attrition or fraud, businesses can trigger automated responses that enhance customer protection and system reliability. The following sections detail the architectural components, the literature governing current trends, and the results achieved using a Python-based streaming implementation environment. This work sets a new standard for intelligent, high-velocity risk management in modern distributed networks.

Literature Survey:

Early research into transaction monitoring focused on statistical methods and rule-based induction applied to static, historical datasets. Verbeke et al. highlighted the importance of building comprehensible prediction models, noting that business stakeholders require transparency as much as accuracy. While these early works established a baseline for identifying behavioral and transactional features, they relied on batch processing, which cannot handle the high velocity of modern data streams or adapt to rapid shifts in consumer behavior.

The shift toward machine learning in financial sectors became more prominent with the exploration of decision trees and support vector machines for usage pattern modeling.



Researchers demonstrated that billing information and transaction frequency could be modeled to predict anomalies with reasonable precision. However, a persistent challenge remains in the form of "concept drift," where the statistical properties of transaction data change over time. Without adaptive mechanisms, the predictive accuracy of these static models inevitably degrades as new fraud tactics emerge.

Nakamoto's introduction of blockchain technology in 2008 marked a major shift toward decentralized trust management. By providing an immutable, transparent, and cryptographic ledger, blockchain addressed the vulnerabilities of centralized authentication servers. Researchers quickly recognized that these principles were directly applicable to securing transaction logs and access control records in distributed environments. Current literature emphasizes that combining blockchain's immutability with automated policy enforcement through smart contracts creates a robust foundation for secure distributed systems.

Recent studies have begun to explore the use of deep learning and Transformer models for more granular data inspection. Traditional machine learning models often rely on handcrafted features, whereas deep learning architectures can capture complex dependencies within network and transaction payloads. Transformer-based models, in particular, utilize self-attention mechanisms to identify subtle indicators of threat that traditional models might miss. These advancements allow for a more holistic

analysis of transaction sequences rather than just individual, isolated events.

Finally, the business value of these technical implementations is contextualized by the need for actionable recommendations. Provost and Fawcett argued that predictions are insufficient without a clear link to strategic action, stressing that risk scores must be translated into tangible interventions. Large Language Models (LLMs) have recently been introduced to provide explainable feedback for detected anomalies, helping developers and managers understand "why" a segment was flagged. This evolution toward intelligent, explainable, and real-time systems represents the current frontier in automated transaction security.

Methodology

The proposed system architecture is designed to handle the velocity and variety of streaming financial data through a modular framework. The methodology integrates blockchain security with machine learning analytics to provide an end-to-end secure transaction environment. The first phase involves the Streaming Data Ingestion Module, which collects interaction data from sources such as mobile applications, transaction logs, and management hubs. This module ensures low-latency capture and prepares the data for the real-time processing pipeline.

Following ingestion, the data enters the Preprocessing and Feature Engineering Module. Unlike traditional batch systems, this module cleans and normalizes data on the fly,



removing noise and handling missing values in the stream. Advanced techniques extract syntactic and structural features, such as "rolling window" transaction frequencies and recent inactivity durations. These extracted features help the machine learning models understand the underlying structure and context of the transaction sequence.

The core analytical engine resides in the Real-Time Anomaly Prediction Module, which utilizes incremental learning algorithms. These models compute risk probability scores for individual transactions by comparing current behavior against learned historical patterns. The models are designed to detect early warning signs of fraud and update their internal parameters dynamically as each new event arrives. This adaptive approach allows the system to mitigate the effects of concept drift and remain effective in dynamic market conditions.

Parallel to the prediction engine is the Blockchain-Based Security and Authorization Module. Every access request and transaction decision is verified through smart contracts deployed on a decentralized ledger. User identities and permissions are securely recorded to ensure integrity and traceability. By decentralizing the authorization process, the framework eliminates single points of failure and provides a tamper-proof audit trail for all system activities. This integration ensures that the anomaly detection insights are coupled with robust enforcement mechanisms.

The final phase is the Reporting, Visualization, and Feedback Module, which presents results

via real-time dashboards. This interface displays live risk levels, anomaly distribution, and individual user risk scores. Visual elements such as charts and highlighted transaction sections improve comprehension for decision-makers. The module also generates explainable feedback, helping security teams understand the logic behind specific alerts. This transforms raw analytical data into practical, actionable insights for proactive financial risk management.

Results and Discussion

The implementation of the proposed framework was conducted using Python in a Google Colab environment, leveraging cloud-based computing resources for scalability. The system was tested using a simulated stream of transaction data to evaluate performance under continuous load. Key metrics monitored during the evaluation included authentication success rates, feature distribution quality, and the accuracy of risk predictions. The experimental setup utilized a hardware configuration with 4GB of RAM and a Dual Core processor to simulate standard enterprise node capabilities.

Analysis of the User Identity Authentication Module showed a steady improvement in reliability over successive attempts. Initial authentication success began at 85% and reached 99% as the decentralized mechanism stabilized. This upward trend indicates that the blockchain-based credentials effectively verified user identities while reducing failures over time. The consistent improvement reflects the framework's



adaptability and its ability to provide secure access without reliance on a central authority.

The Feature Extraction and Preprocessing results were visualized using a histogram of normalized transaction features. The distribution followed a near-Gaussian pattern, confirming that noise and outliers were effectively removed from the raw data streams. High-quality feature extraction is a critical prerequisite for accurate anomaly detection, as it ensures that the models are analyzing relevant behavioral signals. These results validate the robustness of the system's initial data preparation layers.

The Real-Time Anomaly Prediction Module provided a detailed comparison of risk across different system components. For example, the "Database Layer" exhibited the highest bug/risk score, exceeding 0.8, while the "API Layer" showed a lower risk of approximately 0.55. This granular risk assessment helps security teams prioritize their monitoring and response efforts on the most vulnerable parts of the infrastructure. By identifying these high-risk areas in real time, the system allows for immediate intervention to prevent potential financial loss.

Finally, the Feedback and Authorization Distribution showed that the system correctly handled a high volume of requests. Out of a simulated batch, 780 requests were granted access based on low risk, while 220 were denied due to suspicious patterns. The "Audit Log Growth" followed a linear scale, increasing from 100 to 1,150 transactions as block numbers reached 10. These findings collectively prove that the

framework provides a secure, transparent, and scalable solution for real-time anomaly detection in large-scale financial environments.

Conclusion

This research has successfully demonstrated the development of a Real-Time Anomaly Detection Framework for transaction data using machine learning and blockchain technology. By moving away from reactive batch processing, the system provides organizations with the agility to identify fraudulent activities as they happen. The integration of smart contracts ensures that security policies are enforced automatically and consistently, while the immutable ledger provides transparency and accountability for compliance. The experimental results validate that the system maintains high accuracy and scalability even under increasing data loads. Ultimately, this framework offers a proactive and robust solution for protecting digital financial ecosystems against modern cyber threats.

Future Work

In the future, the Real-Time Anomaly Detection system can be enhanced by integrating Decentralized Identity (DID) frameworks to provide users with greater control over their security credentials. This would further improve privacy and reduce dependence on external identity providers. Additionally, incorporating more advanced Deep Learning architectures, such as LSTMs or RNNs, could improve the detection of

complex temporal sequences in transaction behavior. Adopting Layer-2 scaling techniques will be essential for reducing transaction latency as the network scales to global volumes. Finally, the system can be extended to support context-aware access control, utilizing factors like location and device trust levels to provide even more granular risk assessments in real-time environments.

References

1. T. Verbeke, D. Martens, C. Mues, and B. Baesens, "Building comprehensible customer churn prediction models with advanced rule induction techniques," *Expert Systems with Applications*, vol. 38, no. 3, pp. 2354-2364, Mar. 2011.
2. Google, "Google Colab: Cloud-Based Python Environment," Product Guide, 2022.
3. S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," White paper, 2008.
4. K. Zhang, X. Liang, J. Ni, K. Yang, and X. Shen, "Exploiting blockchain for secure and efficient access control in decentralized networks," *IEEE Transactions on Communications*, vol. 66, no. 11, pp. 5646-5659, Nov. 2018.
5. A. Vaswani et al., "Attention is all you need," in *Proc. 31st Int. Conf. Neural Information Processing Systems (NeurIPS)*, 2017.
6. J. Gama, I. Žliobaitė, A. Bifet, M. Pechenizkiy, and A. Bouchachia, "A survey on concept drift adaptation," *ACM Computing Surveys*, vol. 46, no. 4, pp. 1-37, Mar. 2014.
7. M. Allamanis, E. T. Barr, C. Bird, and C. Sutton, "A survey of machine learning for big code and naturalness," *ACM Computing Surveys*, vol. 51, no. 4, pp. 1-37, Aug. 2018.
8. Real-Time Anomaly Detection in Transaction Data Using Machine Learning.pdf, project documentation, 2023.
9. J. Xu, K. Xue, and S. Li, "A blockchain-based access control framework for cloud computing," *IEEE Access*, vol. 7, pp. 1-14, 2019.
10. Y. Zuo, S. Zhang, and Y. He, "Blockchain-based access control for Internet of Things," *IEEE Internet of Things Journal*, vol. 7, no. 6, pp. 1-12, June 2020.
11. M. Ali, J. Nelson, R. Shea, and M. J. Freedman, "Blockstack: A global naming and storage system secured by blockchains," in *Proc. USENIX Annu. Tech. Conf.*, 2016.
12. Microsoft, "Introduction and Core Philosophy of Windows 11," *Windows Technical Documentation*, 2021.
13. Python Software Foundation, "Python History and Key Features," *Python Documentation*, 2023.
14. ISO/IEC, "Unified Modeling Language (UML) Specification," 2017.
15. F. Provost and T. Fawcett, "Data science for business: What you need to know about data mining and data-analytic



INTERNATIONAL JOURNAL OF COMPUTER SCIENCE

Scholarly Peer Reviewed Research Journal - PRESS - OPEN ACCESS

ISSN: 2348-6600



Since 2012

www.ijcsjournal.com

Volume 14, Issue 1, No 25 2026

ISSN: 2348-6600

REFERENCE ID: IJCS-695

PAGE NO: 094-098

thinking," IEEE Intelligent Systems, vol.
28, no. 6, pp. 63-67, 2013.