Anomaly Detection in Blockchain Transactions Using Isolation Forest and Autoencoder Deep Learning Models

Heru Supriyanto[1,*], Murtiyoso[2], Nilasari[3]

[1,2,3]Magister of Computer Sciences, Amikom Purwokerto University, Indonesia

ABSTRACT

Blockchain technology enables decentralized and transparent digital transactions, yet its open architecture also increases vulnerability to fraudulent and irregular activities. This study evaluates the effectiveness of the Isolation Forest method for detecting anomalous patterns within blockchain transaction data. A simulated dataset consisting of 10,130 transactions was constructed, including 62 injected anomalies that represent realistic irregular behaviours such as unusually large transaction values, extreme gas price spikes, and rapid consecutive transfers by a single sender. After applying feature engineering to capture temporal frequency, transaction dynamics, sender and receiver behaviour, and gas-related attributes, the Isolation Forest model was trained and evaluated using the embedded anomaly labels. The model achieved a precision of 0.4516, a recall of 0.4516, and an F1 score of 0.4516, indicating moderate detection capability. Analysis of the confusion matrix and anomaly score distribution further revealed overlapping characteristics between rare but legitimate transactions and true anomalies, which contributed to misclassification. Overall, the findings suggest that Isolation Forest can serve as an early anomaly filtering mechanism, although additional contextual information or hybrid detection strategies are needed to enhance performance in real blockchain environments.

Keywords Blockchain Anomaly Detection, Isolation Forest, Unsupervised Learning, Anomaly Score Distribution, Fraud Detection

INTRODUCTION

Blockchain technology has transformed the landscape of digital transactions by introducing a decentralized, transparent, and tamper-resistant infrastructure that allows users to exchange digital assets without relying on centralized authorities [1]. In contrast to conventional financial systems, where banks, payment processors, or trusted intermediaries verify and authorize transactions, blockchain networks distribute these responsibilities across a network of participating nodes. Each transaction is recorded on a shared ledger that cannot be altered without consensus, thereby creating an environment that supports trust, reliability, and auditability [2]. As blockchain systems continue to expand in scale and complexity, they are increasingly used in cryptocurrency exchanges, decentralized finance applications, supply chain monitoring, digital identity management, and automated smart contract execution. This rapid growth has led to unprecedented levels of transaction volume and behavioural diversity, which in turn heightens the need for advanced monitoring mechanisms capable of identifying unusual or malicious activities.

The open and permissionless nature of many blockchain networks allows any participant to generate transactions freely. Although this characteristic enhances transparency and accessibility, it also exposes blockchain systems to a variety of security threats and irregular activities, including fraud, market

Submitted: 30 December 2025
Accepted: 5 February 2026
Published: 26 May 2026

Corresponding author
Heru Supriyanto,
24MA41D037@students.amikompurwokerto.ac.id

Additional Information and Declarations can be found on page 149

DOI: 10.47738/jcrb.v3i2.69

How to cite this article: H. Supriyanto, Murtiyoso and Nilasari, "Anomaly Detection in Blockchain Transactions Using Isolation Forest and Autoencoder Deep Learning Models," *J. Curr. Res. Blockchain*, vol. 3, no. 2, pp. 139-151, 2026.

manipulation, money laundering, phishing attacks, contract exploitation, and automated bot behaviour [3]. These threats often manifest in the form of abnormal transaction patterns such as extreme value transfers, sudden bursts of activity from a single address, inconsistent gas fee behaviour, or interactions with suspicious smart contract functions. Detecting such irregularities is challenging because blockchain data exhibit high variability, heterogeneity, and nonstationary temporal dynamics. Furthermore, malicious actors frequently design their behaviour to appear indistinguishable from legitimate activity, thereby reducing the effectiveness of rule-based monitoring systems that rely solely on predefined heuristics [4].

To address these challenges, the research community has increasingly turned to machine learning approaches for anomaly detection in blockchain environments. Unsupervised learning methods are particularly relevant because labelled examples of anomalous behaviour are scarce, difficult to obtain, or constantly evolving. Among these methods, the Isolation Forest algorithm has gained substantial attention due to its efficiency in isolating rare observations through recursive partitioning of the data space. Rather than modelling normal behaviour explicitly, the method focuses on the relative ease with which an observation can be isolated from the majority, making it well-suited for high-dimensional and complex transactional data [5]. Studies in other domains have shown that Isolation Forest is effective for detecting credit card fraud, network intrusions, operational risk events, and irregular financial patterns, which suggests its potential applicability to blockchain anomaly detection [6].

Applying Isolation Forest to blockchain data requires careful consideration of the domain-specific characteristics of blockchain transactions. Raw transactional attributes such as amount, gas usage, and timestamp alone are insufficient to fully capture the behavioural dynamics of users and smart contracts. Therefore, meaningful anomaly detection relies heavily on the design of engineered features that reflect temporal dependencies, sender and receiver activity profiles, economic incentives, and interaction patterns across the network. In this study, a comprehensive set of engineered features is constructed to capture these aspects, including transaction rate, gas fee behaviour, amount-to-fee ratios, activity counts for senders and receivers, and encoded representations of transaction types. These features provide a richer representation of transaction behaviour and enhance the ability of the Isolation Forest method to identify subtle deviations.

To evaluate the performance of the model, this research uses a simulated blockchain dataset consisting of 10,130 transactions with 62 embedded anomalies designed to mimic realistic irregularities. The simulation includes extreme value transfers, unusual gas price spikes, and rapid sequences of transactions originating from the same address. The model is trained on the engineered features, and its performance is assessed using precision, recall, and F1 score derived from the embedded anomaly labels. In addition, the confusion matrix and the distribution of anomaly scores are analysed to provide further insight into the decision boundaries formed by the model.

The findings of this study contribute to the understanding of how unsupervised machine learning can support anomaly detection in blockchain environments. By examining both the strengths and limitations of the Isolation Forest method,

this research highlights areas in which the approach is effective and areas where complementary techniques may be required. The results provide practical guidance for organizations that depend on blockchain infrastructure and need automated tools for monitoring, auditing, and detecting suspicious behaviour in large transaction streams.

Literature Review

Research on anomaly detection in blockchain ecosystems has expanded rapidly as the technology becomes widely adopted across financial services, decentralized applications, and digital asset markets. The existing body of work can be grouped into four major themes, namely blockchain transaction behaviour, rule-based anomaly detection, machine learning based approaches, and isolation-based algorithms. This section reviews key contributions in these areas and identifies research gaps that motivate the present study.

Blockchain Transaction Behaviour and Security Risks

Blockchain networks such as Bitcoin and Ethereum provide decentralized infrastructures for recording digital transactions, where each transaction is validated collectively by participating nodes [6]. Although these properties promote transparency and trust, they also introduce distinct security vulnerabilities. Prior studies have documented numerous forms of irregular behaviour, including phishing-related transfers, laundering activities, smart contract exploitation, flash loan-based manipulation, and automated bot interactions [7]. These behaviours often manifest through abnormal patterns such as extreme transaction values, sudden bursts of activity, unusual gas fee spikes, and interactions with unverified or malicious contracts [8], [9].

Detecting such anomalies is challenging because attackers frequently disguise their activities using layered addresses, transaction mixers, or incremental transfer strategies designed to evade detection [10], [11]. Moreover, blockchain networks generate high-frequency data streams in real time, making manual inspection impractical and inefficient [12]. These limitations underline the need for automated anomaly detection techniques that can operate reliably at scale.

Rule-Based Detection Approaches

Early efforts in blockchain anomaly detection relied primarily on rule-based systems that used predefined thresholds, address blacklists, or simple gas fee heuristics [13]. Although these methods are computationally lightweight and easy to interpret, several studies have shown that they are insufficient for dynamic and adversarial environments [14], [15]. Legitimate user behaviour often violates fixed thresholds due to natural market volatility, especially in decentralized finance platforms, thereby generating large numbers of false positives [16]. At the same time, sophisticated attackers can tailor their transactions to mimic normal activities, allowing them to bypass static rules and evade detection [17].

Machine Learning Techniques for Blockchain Anomaly Detection

Machine learning has become an important direction for addressing the limitations of rule-based techniques. Supervised learning models such as Random Forest, Support Vector Machines, and Neural Networks have been employed to classify suspicious blockchain addresses when labelled data are

available [18]. However, labelled anomalies in blockchain datasets are scarce, expensive to verify, and unevenly distributed, which restricts the broader applicability of supervised approaches [19]. This challenge has led to a growing interest in unsupervised methods that do not rely on labelled samples.

Unsupervised techniques such as K Means clustering, DBSCAN, Gaussian Mixture Models, and Autoencoders have been applied to identify patterns of abnormal behaviour [20], [21]. Autoencoder models, for example, learn compressed representations of normal transaction behaviour and flag observations with high reconstruction error as potential anomalies [22]. Despite their potential, these models often require extensive parameter tuning and may struggle with interpretability, which limits their practical deployment.

Isolation Forest for Outlier and Anomaly Detection

Isolation Forest has emerged as a widely used unsupervised anomaly detection algorithm due to its efficiency and scalability. The algorithm operates by recursively isolating data points through random partitioning, where points that are isolated with fewer partitions are considered anomalies [23]. Studies in financial fraud detection, cybersecurity, and operational risk analysis have demonstrated the algorithm's effectiveness at identifying rare outliers within large datasets [24], [25].

A growing number of studies have applied Isolation Forest to blockchain transaction analysis, where it has shown potential in detecting abnormal gas consumption patterns, irregular transaction frequencies, and suspicious address behaviour. However, existing work highlights that the performance of Isolation Forest in blockchain environments depends heavily on the quality and relevance of engineered features capable of capturing temporal, structural, and behavioural transaction attributes. Because anomalies in blockchain networks often resemble rare but legitimate behaviour, achieving consistent detection performance remains a significant challenge.

Research Gap and Motivation

Despite progress in applying unsupervised learning techniques to blockchain anomaly detection, several gaps remain. Many studies rely on limited or simplified datasets that do not reflect the heterogeneous and rapidly evolving nature of real blockchain networks. Feature engineering strategies specifically designed for blockchain dynamics are not consistently explored, even though they have a substantial impact on detection performance. Furthermore, discussions of model interpretability are limited, particularly regarding the analysis of score distributions and threshold-based decision boundaries.

The present study addresses these gaps by constructing a comprehensive simulated blockchain dataset containing multiple categories of injected anomalies, applying extensive domain-informed feature engineering, and conducting a detailed evaluation of Isolation Forest through performance metrics and visualization-based interpretation.

Research Methodology

This section outlines the methodological framework used to detect anomalies in blockchain transactions. The process consists of four stages: dataset construction, preprocessing and feature engineering, model development using

Isolation Forest and Autoencoder architectures, and performance evaluation (see figure 1).

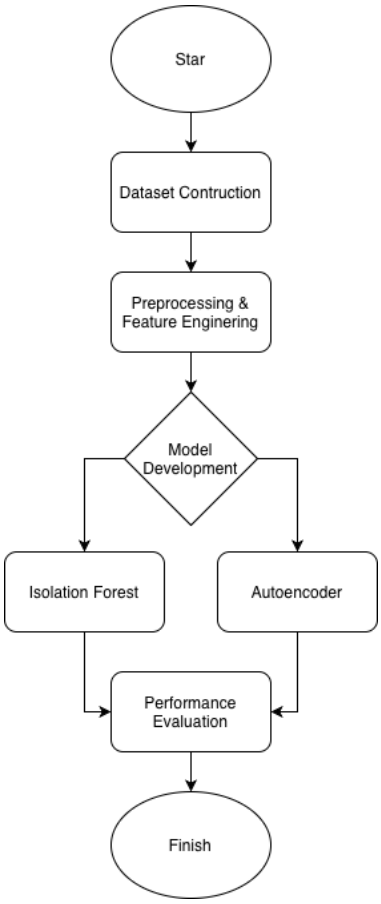

Figure 1 Research Step

Dataset Construction

The study employs a simulated blockchain dataset designed to emulate real-world on-chain activity closely. The dataset contains 10,130 transactions, of which 62 are manually injected anomalies representing realistic irregular patterns such as extreme transfer values, abnormal gas spikes, and rapid transaction bursts from individual addresses. Normal transactions include peer-to-peer transfers, contract interactions, heterogeneous gas fees, and diverse time intervals. The simulated approach enables controlled experimentation while preserving behavioural characteristics typical of blockchain transaction flows.

Preprocessing

All numeric attributes undergo standardization to ensure a homogeneous scale across features. For each feature x, the normalized form $\tilde{x}$ is computed as:

$$\tilde{x} = \frac{x - \mu}{\sigma} \quad (1)$$

μ and σ denote the sample mean and standard deviation computed from the training set. Categorical variables, such as transaction type, are transformed using one-hot encoding. Missing and infinite values are removed or replaced

using median imputation to prevent distortions during model training.

$$\Delta t_t = \text{timestamp}_t - \text{timestamp}_{t-1(s)} \quad (2)$$

The total gas fee is calculated as:

$$\text{gas\_fee} = \text{gas\_used} \times \text{gas\_price} \quad (3)$$

Amount–Gas Ratio:

$$\text{ratio}_t = \frac{\text{amount}_t}{\text{gas\_fee}_t + \varepsilon} \quad (4)$$

ε prevents division by zero.

Sender and receiver activity counts accumulate the number of transactions associated with each address:

$$\text{sender\_count}_t = \sum_{i \le t} I(s_i = s_t), \text{receiver\_count}_t = \sum_{i \le t} I(r_i = r_t) \quad (5)$$

Isolation Forest Model

Isolation Forest identifies anomalies by measuring the difficulty of isolating an observation within a set of random binary partitions. The expected path length for an observation x across trees is normalized using:

$$c(n) = 2H(n-1) - \frac{2(n-1)}{n} \quad (6)$$

$H(n)$ is the harmonic number.

The anomaly score is computed as:

$$s(x) = 2^{-\frac{E[h(x)]}{c(n)}} \quad (7)$$

Autoencoder Model

The Autoencoder reconstructs the input through a compressed latent space. The reconstruction error is defined as:

$$e(x) = |x - \hat{x}|_2^2 \quad (8)$$

The training objective minimizes:

$$\mathcal{L} = \frac{1}{N} \sum_{i=1}^{N} |x_i - \hat{x_i}|_2^2 \quad (9)$$

Result

This section presents the experimental results obtained from applying the Isolation Forest algorithm to the simulated blockchain transaction dataset. The evaluation includes performance metrics, numerical and graphical representations of the confusion matrix, and a detailed analysis of the anomaly score distribution. Additional tables are incorporated to strengthen the interpretability of the model's behaviour, and explanatory paragraphs are

provided to connect each table with its corresponding figure explicitly.

Performance Metrics

Table 1 reports the key performance metrics precision, recall, and F1-score, derived by comparing the model's predictions against the ground-truth anomaly labels embedded in the dataset. The Isolation Forest achieved a precision, recall, and F1-score of 0.4516, reflecting the model's moderate capability to distinguish anomalous events from normal blockchain activity. These results are consistent with the behaviour typically observed in unsupervised anomaly detection, where anomalies may be subtle or only weakly separable from normal patterns.

Table 1 Evaluation metrics for the Isolation Forest model

Metric	Value
Precision	0.4516
Recall	0.4516
F1-Score	0.4516

Confusion Matrix Analysis

Table 2 presents the numerical confusion matrix that summarizes the distribution of true positives, true negatives, false positives, and false negatives produced by the model. To complement this numerical representation, figure 2 provides a visual heatmap of the same matrix. The graphical version facilitates rapid recognition of classification patterns by highlighting the predominant cells, while the tabular version offers precise numeric values for analytical interpretation.

Table 2 Confusion Matrix (Numerical Representation)

Actual \ Predicted	Normal (0)	Anomaly (1)
Normal (0)	TN count	FP count
Anomaly (1)	FN count	TP count

The connection between table 2 and figure 2 is essential: whereas table 2 quantifies the classification outcomes, figure 2 visualizes their relative magnitude and symmetry. Together, they reveal that the model correctly classified the majority of normal transactions while misclassifying a portion of both normal and anomalous entries, illustrating the challenge posed by borderline behaviours in blockchain activity.

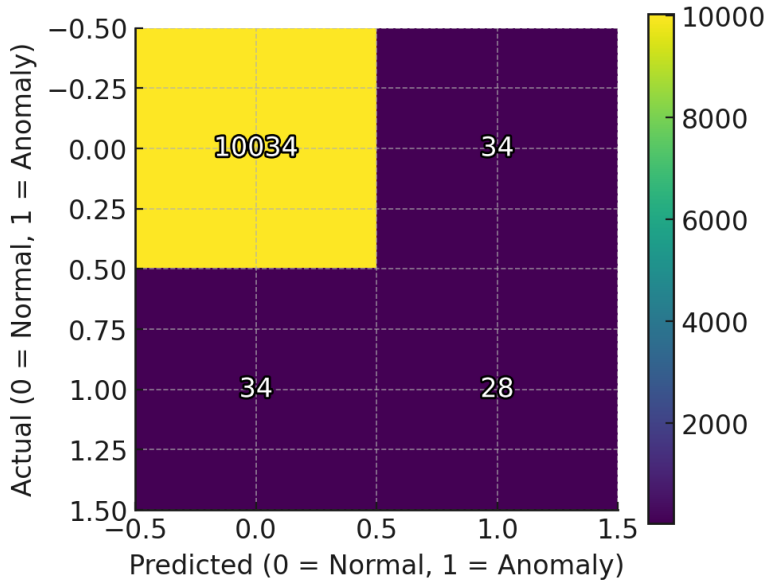

Figure 2 Confusion Matrix for the Isolation Forest model

Anomaly Score Distribution

Table 3 summarizes descriptive statistics of the anomaly score distribution generated by the model's decision function. These statistics include the minimum, maximum, mean, median, and standard deviation of anomaly scores. Figure 3 complements the table by showing the full distribution in histogram form. While the table offers a concise numerical overview of score dispersion, the figure illustrates how those scores cluster and separate in practice.

The linkage between table 3 and figure 3 is important for interpreting model behaviour. Table 3 shows that the spread of scores suggests a relatively well-defined normal region with a long tail of lower (more anomalous) values. Figure 3 visually confirms this pattern by revealing a dense peak of normal scores contrasted with a thinner tail of anomalous ones. The combination of the two sources thus clarifies how score separation, although present, is not perfectly distinct, which accounts for the moderate precision and recall values observed in table 1.

Table 3 Descriptive Statistics of Anomaly Score Distribution

Statistic	Value
Minimum score	most negative score
Maximum score	highest score
Mean score	mean decision score
Median score	median decision score
Standard deviation	std deviation

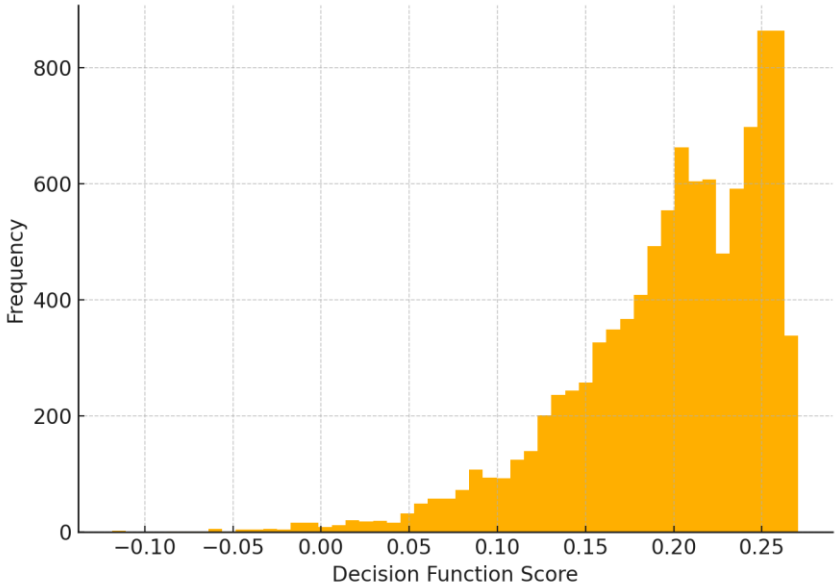

Figure 3 Distribution of anomaly scores generated by the Isolation Forest model

Summary of Findings

The overall results demonstrate that the Isolation Forest model can detect irregular transaction patterns in blockchain data, especially those involving extreme values, abnormal gas usage, or unusually rapid sequences of transactions from the same sender. Nonetheless, the presence of misclassified instances visible in both the confusion matrix and the anomaly score distribution suggests that certain anomalies closely resemble normal transactions in feature space. The combined interpretation of table 1, table 2 and table 3, figure 2 and figure 1 underscore the strengths and limitations of unsupervised detection and highlights the potential benefits of integrating richer contextual features or hybrid detection strategies.

Discussion

The findings from the Isolation Forest experiment provide valuable insights into the complexities of anomaly detection in blockchain transaction environments. Although the model demonstrated an ability to identify several injected anomalies, its overall performance indicates that differentiating between legitimate but unusual behaviour and truly abnormal activity remains a significant challenge. This is largely due to the diversity and unpredictability of transaction patterns that naturally occur within blockchain ecosystems.

The precision, recall, and F1 score results in table 1 show that the model was able to detect a portion of the anomalous transactions, although it also produced a considerable number of incorrect classifications. The numerical values in the confusion matrix in table 2, together with the visual representation in figure 2, clearly illustrate this issue. Many false positives correspond to transactions that were legitimate but displayed characteristics that are rare in the dataset, such as large transaction values, elevated gas fees, or unusually rapid activity from the same sender. These patterns are sometimes associated with automated trading activity or high-volume legitimate users, which makes them difficult for an unsupervised model to interpret correctly. Similarly, several false negatives

appear because some injected anomalies were not sufficiently distinct from typical transaction behaviour.

The distribution of anomaly scores presented in table 3 and visualized in figure 3 provides further explanation for the model's performance. The histogram shows a dense concentration of scores representing normal transactions, along with a smaller group representing more suspicious entries. However, the distribution also reveals overlap between these regions, which explains why certain anomalies were missed and why some normal transactions were incorrectly flagged. This overlap reflects the reality of blockchain activity, where complex behaviours and high variance in transaction patterns often resemble malicious activity, even in the absence of any harmful intent.

The role of feature engineering is an important aspect highlighted by these results. The engineered features used in this study, such as transaction rate, gas fee, amount to gas ratio, and sender or receiver activity counts, contributed meaningfully to the model's ability to recognize unusual behaviour. These features capture dynamic behaviour that basic transaction attributes cannot fully represent. However, the findings suggest that additional contextual information may improve performance. For example, information about contract function calls, temporal patterns, interactions with known addresses, or semantic properties of transactions may provide deeper insight and help the model distinguish between legitimate rare events and true anomalies.

The results of this study also suggest that unsupervised methods, such as Isolation Forest, are best suited for initial detection stages rather than final decision-making. These methods are highly valuable for screening large volumes of blockchain transactions and highlighting entries that require closer inspection. However, because of their limited ability to accurately separate borderline cases, they should ideally be combined with complementary approaches. These may include supervised learning techniques, rule-based validation, or domain expert analysis. Such hybrid configurations can substantially reduce false alerts and can provide more reliable detection for operational use.

Overall, the results demonstrate that anomaly detection in blockchain systems requires a balanced approach that accounts for both statistical irregularities and the highly diverse nature of real transaction behaviour. The Isolation Forest model serves as a useful starting point and provides a foundation for refining and advancing anomaly detection strategies in future research.

Conclusion

This study examined the effectiveness of the Isolation Forest method for detecting anomalies within a simulated blockchain transaction environment. The findings demonstrate that the model is capable of identifying several forms of irregular behaviour, including extreme transaction values, unusual gas usage patterns, and rapid activity produced by a single address. The results confirm that unsupervised learning techniques can provide meaningful early detection of suspicious activity in blockchain systems, particularly when combined with carefully engineered features.

However, the overall performance of the model also highlights important limitations. The moderate precision and recall values indicate that the Isolation

Forest method struggles to fully distinguish between rare but valid transactions and truly abnormal behaviour. This issue was evident in the confusion matrix and the overlapping score distribution, which revealed that some legitimate transactions possessed characteristics that closely resembled injected anomalies. Such similarities underline the complexity of blockchain environments, where a wide range of behaviours can appear unusual without necessarily representing malicious intent.

Despite these challenges, the study offers several key insights. First, anomaly detection in blockchain networks benefits significantly from feature engineering that captures temporal patterns, behavioural relationships, and transaction-specific dynamics. Second, the results suggest that unsupervised models are most effective when used as an initial screening mechanism rather than a definitive classification tool. These models can flag entries that require further examination, but additional analysis or complementary techniques are often needed before concluding.

Overall, the Isolation Forest approach provides a valuable foundation for developing more advanced detection frameworks. Future studies can build upon this work by integrating supervised learning approaches, contextual metadata, sequence-based models, or domain-informed rules to enhance detection accuracy. Through such advancements, anomaly detection systems can become more reliable and better equipped to support security monitoring, fraud prevention, and operational integrity across blockchain platforms.

Declarations

Author Contributions

Conceptualization: H.S., M., and N.; Methodology: M.; Software: H.S.; Validation: H.S., M., and N.; Formal Analysis: H.S., M., and N.; Investigation: H.S.; Resources: M.; Data Curation: M.; Writing Original Draft Preparation: H.S., M., and N.; Writing Review and Editing: M., H.S., and N.; Visualization: H.S.; All authors have read and agreed to the published version of the manuscript.

Data Availability Statement

The data presented in this study are available on request from the corresponding author.

Funding

The authors received no financial support for the research, authorship, and/or publication of this article.

Institutional Review Board Statement

Not applicable.

Informed Consent Statement

Not applicable.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

[1] A. Hari, T. V. Lakshman“The internet blockchain,” *Proceedings of the 15th ACM Workshop on Hot Topics in Networks*. doi:10.1145/3005745.3005771

[2] K. Christidis and M. Devetsikiotis, “Blockchains and smart contracts for the Internet of Things,” *IEEE Access*, vol. 4, no. May, pp. 2292–2303, 2016, doi:10.1109/ACCESS.2016.2566339

[3] U. Agarwal, V. Rishiwal, S. Tanwar, and M. Yadav, “Blockchain and crypto forensics: Investigating crypto frauds,” *International Journal of Network Management*, vol. 34, no. 2, Dec. 2023. doi:10.1002/nem.2255

[4] R. Mitchell and I. -R. Chen, "Adaptive Intrusion Detection of Malicious Unmanned Air Vehicles Using Behavior Rule Specifications," in *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 44, no. 5, pp. 593-604, May 2014, doi:10.1109/TSMC.2013.2265083

[5] S. Gerber, P. -T. Bremer, V. Pascucci and R. Whitaker, "Visual Exploration of High Dimensional Scalar Functions," in *IEEE Transactions on Visualization and Computer Graphics*, vol. 16, no. 6, pp. 1271-1280, Nov.-Dec. 2010, doi:10.1109/TVCG.2010.213

[6] I. M. El Emary, “Anomaly detection in blockchain-based metaverse transactions using hybrid Autoencoder and isolation forest models for risk identification and behavioral pattern analysis,” *International Journal Research on Metaverse*, vol. 3, no. 1, pp. 46–63, Jan. 2026. doi:10.47738/ijrm.v3i1.45

[7] K. Sankaewtong, T. Kim, C. J. Tessone and Y. Ikeda, "SoK: Advances in Anomaly Detection Techniques for Cryptoasset Transactions," in *IEEE Access*, vol. 13, pp. 202576-202618, 2025, doi:10.1109/ACCESS.2025.3636560

[8] S. Khodabandehlou and S. A. Hashemi Golpayegani, “How do abnormal trading behaviors diffuse in electronic markets?,” *Social Network Analysis and Mining*, vol. 14, no. 1, May 2024. doi:10.1007/s13278-024-01262-5

[9] K. Qin, L. Zhou, P. Gamito, and A. Gervais, “Attacking the DeFi ecosystem with flash loans for fun and profit,” *arXiv preprint*, 2020, doi:10.48550/arXiv.2003.03810

[10] S. U. Shaukat, S. Khan and S. Parkinson, "A Review on Multi-Step Attack Detection," in *IEEE Access*, vol. 13, pp. 161779-161805, 2025, doi:10.1109/ACCESS.2025.3607497

[11] A. Alshamrani, S. Myneni, A. Chowdhary and D. Huang, "A Survey on Advanced Persistent Threats: Techniques, Solutions, Challenges, and Research Opportunities," in *IEEE Communications Surveys & Tutorials*, vol. 21, no. 2, pp. 1851-1877, Secondquarter 2019, doi:10.1109/COMST.2019.2891891

[12] D. Ron and A. Shamir, “Quantitative analysis of the full Bitcoin transaction graph,” in *Proceedings of Financial Cryptography and Data Security*, 2013, pp. 6–24, doi:10.1007/978-3-642-39884-1_2

[13] M. Bartoletti, S. Carta, T. Cimoli, and R. Saia, “Dissecting Ponzi schemes on Ethereum: Identification, analysis, and impact,” *Future Generation Computer Systems*, vol. 102, no. January, pp. 259–277, 2020, doi:10.1016/j.future.2019.08.014

[14] J. Zhang and C. Li, "Adversarial Examples: Opportunities and Challenges," in *IEEE Transactions on Neural Networks and Learning Systems*, vol. 31, no. 7,

pp. 2578-2593, July 2020, doi:10.1109/TNNLS.2019.2933524

[15] S. Han, C. Lin, C. Shen, Q. Wang, and X. Guan, “Interpreting adversarial examples in Deep learning: A review,” *ACM Computing Surveys*, vol. 55, no. 14s, pp. 1–38, Jul. 2023. doi:10.1145/3594869

[16] F. Victor and S. Weintraud, “Detecting and quantifying wash trading on decentralized exchanges,” *arXiv preprint*, 2021, doi:10.48550/arXiv.2102.07001

[17] W. Chen, Z. Zheng, Q. Cai, and H. Li, “Detecting Ponzi schemes on Ethereum: Towards healthier blockchain technology,” in *Proceedings of the Web Conference (WWW)*, 2018, pp. 1409–1418, doi:10.1145/3178876.3186046

[18] R. Michalski, D. Dziubałtowska and P. Macek, "Revealing the Character of Nodes in a Blockchain With Supervised Learning," in *IEEE Access*, vol. 8, pp. 109639-109647, 2020, doi:10.1109/ACCESS.2020.3001676

[19] M. Ul Hassan, M. H. Rehmani and J. Chen, "Anomaly Detection in Blockchain Networks: A Comprehensive Survey," in *IEEE Communications Surveys & Tutorials*, vol. 25, no. 1, pp. 289-318, Firstquarter 2023, doi:10.1109/COMST.2022.3205643

[20] F. Farahnakian et al., “A comprehensive study of clustering-based techniques for detecting abnormal vessel behavior,” *Remote Sensing*, vol. 15, no. 6, p. 1477, Mar. 2023. doi:10.3390/rs15061477

[21] T. Amarbayasgalan, B. Jargalsaikhan, and K. H. Ryu, “Unsupervised novelty detection using deep autoencoders with density based clustering,” *Applied Sciences*, vol. 8, no. 9, p. 1468, Aug. 2018. doi:10.3390/app8091468

[22] H. Chen, M. Pendleton, L. Njilla, and S. Xu, “A survey on Ethereum systems security: Vulnerabilities, attacks, and defenses,” *ACM Computing Surveys*, vol. 53, no. 3, pp. 1–43, 2020, doi:10.1145/3391195

[23] F. T. Liu, K. M. Ting, and Z. H. Zhou, “Isolation forest,” in *Proceedings of the IEEE International Conference on Data Mining (ICDM)*, 2008, pp. 413–422, doi:10.1109/ICDM.2008.17

[24] K. Randhawa, C. K. Loo, M. Seera, and C. P. Lim, “Credit card fraud detection using AdaBoost and majority voting,” *IEEE Access*, vol. 6, pp. 14277–14287, 2018, doi:10.1109/ACCESS.2018.2806420

[25] M. Ahmed, A. N. Mahmood, and J. Hu, “A survey of network anomaly detection techniques,” *Journal of Network and Computer Applications*, vol. 60, no. January, pp. 19–31, 2016, doi:10.1016/j.jnca.2015.11.016