

HYBRID TEMPORAL-GRAPH MACHINE LEARNING FOR ANOMALY DETECTION IN METAVERSE FINANCIAL TRANSACTIONS

Jacques Ishimwe^{1*}, Jonathan Ngugi¹, Mungwarakarama Irene¹

Faculty of Computing and Information Sciences, University of Lay Adventists of Kigali (UNILAK), Rwanda

Email of Corresponding Author: ngugijonathan4@gmail.com

Abstract

Metaverse financial ecosystems generate large-scale, pseudonymous, and rapidly evolving transaction networks, creating significant challenges for fraud detection. Traditional rule-based approaches are insufficient to capture the complex structural, temporal, and behavioral patterns present in these systems. This study proposes a Hybrid Temporal-Graph Anomaly Detection (HTGAD) framework that combines three complementary models: a tabular XGBoost classifier for transaction-level features, a GraphSAGE-MLP model for relational patterns in sender-receiver networks, and an LSTM model for sequential behavioral analysis. These outputs are integrated using a logistic regression meta-layer to produce a unified risk score, while SHAP explainability enhances interpretability. An empirical evaluation on the Metaverse Financial Transactions Dataset (78,600 records sourced from Kaggle) shows that XGBoost alone achieves near-perfect classification, LSTM captures temporal dependencies effectively, and GraphSAGE provides structural insights despite its limited standalone performance. The HTGAD fusion model outperforms all individual components, achieving an AUC-ROC of 1.000, an F1-score of 0.9998, a precision of 0.9996, and a recall of 1.000 on the test set. Feature importance analysis identifies transaction type, transaction amount, temporal activity, and user behavior as the most influential predictors of anomalies. These results demonstrate that integrating tabular, temporal, and graph-based perspectives produces a robust, accurate, and interpretable system for detecting anomalous and potentially fraudulent activity in complex, decentralized Metaverse financial environments.

Keywords: *Metaverse transactions, hybrid anomaly detection, Graph Neural Networks, temporal modeling, XGBoost, explainable AI (SHAP).*

1. INTRODUCTION

The rapid growth of digital technology has changed how financial activities are conducted worldwide. Within Metaverse environments, users engage in financial activities such as buying virtual land, trading digital items, transferring virtual currencies, and paying for services, forming virtual economies that closely resemble real-world financial systems but operate entirely in digital space (Henrique, 2023). As Metaverse platforms continue to grow, the number and complexity of financial transactions increase rapidly, and because users are often identified only by digital usernames or virtual identities, Metaverse financial systems have become vulnerable to fraudulent activities such as phishing attacks, fake investment schemes, market manipulation, and other forms of financial abuse (Aysan et al., 2024).

Transactions occur continuously and involve many users and digital assets, creating complex and constantly changing patterns that are difficult for humans to monitor manually. Studies on large-scale digital transaction systems show that manual inspection or simple monitoring

techniques are no longer effective for handling such complex and dynamic data (Huynh-The et al., 2023). Previous fraud detection relied mainly on rule-based systems and manual checks with fixed thresholds, which are not suitable for modern virtual environments where transaction behavior changes rapidly and varies widely across users (Macak et al., 2022).

To overcome these challenges, machine learning and deep learning techniques have become increasingly important for analyzing digital financial transaction data. Machine learning models such as Random Forest and XGBoost can effectively identify unusual and suspicious transactions (Hasan et al., 2024; Pocher et al., 2023). Deep learning methods further improve detection by capturing complex patterns in transaction sequences and user behavior, including multilayer perceptrons, convolutional neural networks, and long short-term memory networks (Airlangga, 2024). Reviews of existing studies confirm that hybrid and ensemble machine learning approaches that combine different models and feature types provide more reliable results in identifying fraudulent activities (Umer et al., 2023).

Despite substantial progress, existing approaches remain limited because they often focus on a single perspective of the data and fail to integrate complementary views into a unified framework (Prakash, 2026). Such models struggle to handle nonlinear behavior, class imbalance, and evolving transaction patterns common in Metaverse environments, and offer limited interpretability for practical decision-making (Li, 2024). This study therefore proposes and evaluates a Hybrid Temporal-Graph Anomaly Detection (HTGAD) framework that simultaneously analyzes transaction attributes, interaction relationships, and temporal behavior to effectively detect anomalous financial activities in Metaverse transaction data. The specific objectives are: (i) to develop individual machine learning models for detecting anomalous patterns; (ii) to integrate these models into a hybrid framework and evaluate their combined effectiveness; and (iii) to compare the performance of the hybrid model with the individual models using appropriate evaluation metrics.

2. LITERATURE REVIEW

Chen et al. (2025) introduced the Multi-Distance Spatial-Temporal Graph Neural Network (MDST-GNN) to detect anomalies in metaverse transactions, addressing challenges of complex multi-hop dependencies, anonymized time data, and model adaptability. Tested on the Elliptic Bitcoin dataset of over 200,000 transactions, the model outperformed baselines by 1.5% AUC-ROC and 2.9% AUC-PR, proving robust to temporal anonymization and effective at capturing three-hop neighborhood patterns.

Yang et al. (2024) introduced GeneralDyG, a method for detecting anomalies in dynamic graphs. The approach uses a sampling technique to capture local graph structures, a graph-based extractor to combine node and edge information, and a sequence model to account for temporal changes. It was tested on four datasets and outperformed nine to ten other models, achieving up to a 3.6% improvement in accuracy on Bitcoin datasets and an F1 score of 85.19% on SWaT. Kim et al. (2024) studied financial anomaly detection in dynamic networks, testing Temporal Graph Network (TGN) models and comparing them with baseline models. TGN-Mean achieved an AUC of 0.7747, improving over the best baseline by 13%.

The study suggests that future research could explore hybrid models that combine TGN and hypergraph approaches. Jumaní and Raza (2025) studied anomaly detection in blockchain networks, testing five models on a Bitcoin dataset with 5% synthetic anomalies. Supervised tree-based models (Random Forest and XGBoost) achieved the highest accuracy, while unsupervised models performed poorly. The study recommended future research on hybrid models, real-time detection, and improved feature extraction. Cholevas et al. (2024)

systematically reviewed unsupervised learning methods for anomaly detection in blockchain networks, noting that graph-based representations were most effective for capturing transactional relationships. Single unsupervised methods often struggled with complex anomalies, while combinations of unsupervised methods or hybrid approaches improved detection performance. El Emary (2026) proposed a hybrid autoencoder and Isolation Forest model for anomaly detection in blockchain-based Metaverse transactions, reporting an AUC-ROC of 0.952, an F1-score of 0.870, a precision of 0.880, and a recall of 0.860. Prakash (2026) combined Isolation Forest, Autoencoder, and XGBoost in an ensemble, achieving an AUC-ROC of 0.912 and an F1-score of 0.885, though recall of 0.841 indicated that some anomalies were missed.

Despite substantial progress, critical gaps remain that limit the effectiveness of existing approaches in Metaverse financial environments. First, existing studies tend to treat transaction attributes, interaction structures, and temporal behavior as separate problems, relying on single-perspective models that fail to capture the multifaceted nature of anomalous activity. Second, most models are validated on cryptocurrency datasets that do not capture the behavioral richness of Metaverse transactions, in which financial risk is shaped not only by transaction values but also by user engagement patterns such as session duration, login frequency, and evolving interaction behavior. Third, an interpretability gap persists, as many advanced graph and temporal models operate as black boxes, limiting analysts' ability to understand, trust, or act upon anomaly detection results. Finally, there is a clear evaluation gap, with limited empirical evidence demonstrating whether hybrid models truly outperform individual approaches when tested consistently on behavior-rich Metaverse transaction data. These gaps directly motivate the present study.

3. METHODOLOGY

3.1 Research Design

This study adopted a quantitative research design, using secondary data from Kaggle, specifically the Metaverse Financial Transactions Dataset, which contains 78,600 metaverse-based transaction records. The research proceeded in three main stages: (i) development of separate models for tabular, graph, and temporal features; (ii) integration of these models into the HTGAD framework using a logistic regression meta-layer; and (iii) comparative evaluation of the hybrid model against individual models using standard classification metrics.

3.2 Dataset

The Metaverse Financial Transactions Dataset consists of records documented in `metaverse_transactions_dataset.csv` that capture interactions between blockchain wallet addresses within a simulated decentralized financial ecosystem. Numerical features include transaction amount, time of day, login frequency, session duration, and risk score. Categorical features include transaction type, region, age group, and wallet addresses. The target variable classifies transactions into low, moderate, and high-risk groups. For modeling purposes, moderate- and high-risk transactions were combined into a single binary anomalous class, while low-risk transactions were treated as normal behavior. To maintain temporal integrity and avert information leakage, the dataset was chronologically divided into training (70%), validation (15%), and testing (15%) subsets, yielding 55,020 training samples, 11,790 validation samples, and 11,790 test samples, each with 20 features. The training set contained 44,365 normal and 10,655 anomalous transactions, while the test set comprised 9,575 normal and 2,215 anomalous transactions. Data preprocessing involved removing missing values, encoding categorical features, and normalizing numerical variables to standard scales. For

graph-based modeling, relationships between sender and receiver addresses were represented as nodes and edges. For temporal modeling, transaction sequences were constructed for each user to reflect the order and timing of activities. All experiments were conducted using Python with Pandas, NumPy, Scikit-learn, XGBoost, PyTorch Geometric, DGL, Keras, and SHAP, on a Google Colab instance equipped with an Intel Core i5 processor and 16 GB of RAM.

3.3 Model Development

Three independent models were developed, each capturing a distinct representation of the transaction data:

Tabular Model (XGBoost): The XGBoost classifier was trained on transaction-level and behavioral features, including amount, type, session duration, and login frequency. It detects anomalies by learning patterns in individual transactions, effectively distinguishing between normal and suspicious activity (Xu et al., 2023).

Graph Model (GraphSAGE + MLP): GraphSAGE was used to generate node embeddings for users and their transaction interactions. Inductive learning allows the model to generalize to new and previously unseen users, an important property in dynamic Metaverse platforms where participants change frequently (Ferretti et al., 2025). An MLP further processed the learned embeddings to estimate anomaly likelihoods.

Temporal Model (LSTM): The LSTM model was trained on 50,376 sequences, each with a temporal window of five timesteps and nine features. It analyzes temporal patterns such as irregular transaction timing, bursts of activity, or unusual sequences, which may indicate fraudulent behavior (Sebt et al., 2021). The HTGAD framework integrates the three base models through a logistic regression meta-layer. Each base model produces an anomaly probability score based on its specific perspective; these scores are then fused by the logistic regression layer to generate a final unified risk score for each transaction. The meta-learner assigns weights of approximately 13.47% to XGBoost, -0.77% to GraphSAGE, and -0.22% to LSTM, reflecting the complementary contributions of each model. The fusion layer completed training in only 0.02 seconds, demonstrating that combining multiple models does not introduce significant computational overhead. SHAP was applied to the XGBoost component to provide feature-level explanations of anomaly predictions.

3.4 Evaluation Metrics

Model performance was assessed using accuracy, precision, recall, F1-score, AUC-ROC, and AUC-PR. Accuracy measured the overall proportion of correctly classified transactions. Precision quantified the proportion of transactions labeled as anomalous that were truly anomalous. Recall measured the proportion of true anomalous transactions successfully detected. The F1-score combined precision and recall into a single harmonic mean. The AUC-ROC summarized the model's ability to discriminate between normal and anomalous transactions across all possible thresholds, while the AUC-PR was particularly relevant for the imbalanced class distribution in this dataset.

4. RESULTS

4.1 Dataset Exploration

The dataset exhibited a clear class imbalance: 80.8% of cases (approximately 63,552 observations) belonged to the low-risk group, while the high-risk category comprised 11.0% (approximately 8,646 cases) and the moderate-risk group 8.3% (approximately 6,526 observations). For modeling, moderate- and high-risk transactions were combined into a single anomalous class, resulting in approximately 19.2% anomalous and 80.8% normal transactions.




Figure 1: Distribution of Observations by Risk Category

Transaction amount analysis revealed that moderate-risk transactions tended to have higher median amounts and wider spreads than low- and high-risk transactions. Anomalous transactions were generally concentrated at higher transaction amounts, while normal transactions were more common in the mid-to-lower range, although overlap between the two distributions indicated that transaction amount alone is not sufficient to fully distinguish anomalous behavior. Monthly temporal analysis showed that total transaction counts remained relatively stable throughout 2022, with anomaly rates remaining within a narrow range of approximately 0.18 to 0.20. Analysis of transaction type and region revealed that Mobile transactions exhibited the highest anomaly rate at approximately 0.65, followed by Online at 0.55, while In-store transactions displayed the lowest rate near 0.15. Geographically, Africa recorded an elevated anomaly rate of approximately 0.6, in contrast to North America and Asia, which showed rates below 0.2. Session behavior analysis revealed that anomalous users appeared more frequently at 6 to 8 logins than at 3 to 5 logins for normal users, with wider session duration variation and higher risk scores. The Pearson correlation matrix indicated a strong positive association between login_frequency and session_duration ($r = 0.872$) and between risk_score and the binary label ($r = 0.817$), confirming that risk_score is a significant predictor of anomalous behavior. Veteran users exhibited the highest anomaly rate at approximately 0.31, followed by new users at about 0.26, while established users showed a near-zero anomaly rate.

4.2 Tabular Model (XGBoost) Performance

The XGBoost classifier demonstrated near-perfect performance across all dataset splits. On the training set, the model achieved an AUC-ROC of 100.00%, AUC-PR of 99.99%, and overall accuracy of 100.00%. For anomalous transactions, precision was 99.98%, recall was 100.00%, and the F1-score was 99.99%. On the validation set, accuracy remained at 99.99%, precision for anomalies was 99.96%, and recall reached 100.00%. On the test set, which represented completely unseen data, the model continued to perform exceptionally, achieving AUC-ROC of 100.00%, AUC-PR of 100.00%, accuracy of 99.99%, precision of 99.95%, recall of 100.00%, and F1-score of 99.98%. These results demonstrate strong generalization.




Figure 2: Performance Comparison of the XGBoost Model Across Training, Validation, and Test Datasets

4.3 Graph Model (GraphSAGE + MLP) Performance

The GraphSAGE component was trained for 0.8 seconds, while the downstream MLP classifier required 12.7 seconds over 12 iterations. The best validation score was 0.8064. On training data, overall accuracy was 80.63%, with AUC-ROC of 50.95% and AUC-PR of 19.63%. Critically, the precision, recall, and F1-score for anomalous nodes were all 0.0%, indicating that the model classified nearly all nodes as normal. Validation and test performance followed a similar pattern, with accuracy of 81.03% and 81.21% respectively, while class-specific metrics remained at 0.0%. These results confirm that degree-only node features are insufficient for anomaly detection; however, the structural embeddings still provide complementary relational information that enhances the HTGAD ensemble. The binary cross-entropy loss began at approximately 0.175 in the first epoch and gradually decreased to approximately 0.167 by epoch 20. Node-level accuracy started near 0.46 and steadily rose to about 0.53 by epoch 20, indicating the model was learning to distinguish anomalies rather than defaulting to the majority class.




Figure 3: GraphSAGE Model Evaluation Across Dataset Splits with Zero Values Marked

4.4 Temporal Model (LSTM) Performance

The LSTM model was trained over ten epochs, with training loss decreasing from 0.8494 to approximately 0.1782 and validation loss from 0.3611 to approximately 0.0971, indicating effective learning of temporal patterns. Training was completed in 251 seconds. On the training set, LSTM achieved an AUC-ROC of 99.87%, AUC-PR of 99.49%, accuracy of 98.15%,

precision of 92.79%, recall of 98.08%, and F1-score of 95.36%. Performance on validation and test sets remained consistently high, with test AUC-ROC at 99.89%, AUC-PR of 99.53%, accuracy of 98.41%, precision of 93.86%, recall of 98.07%, and F1-score of 95.92%, confirming strong generalization and minimal overfitting.




Figure 4: Performance Comparison of the LSTM Model Across Training, Validation, and Test Datasets

4.5 HTGAD Fusion Model Performance

The HTGAD hybrid fusion framework achieved near-perfect classification across all splits. On the training set, AUC-ROC and AUC-PR were both 100.00%, accuracy was 100.00%, precision for anomalous transactions was 99.98%, recall was 100.00%, and F1-score was 99.99%. Validation results showed AUC-ROC of 100.00%, AUC-PR of 99.99%, accuracy of 99.99%, precision of 99.96%, recall of 100.00%, and F1-score of 99.98%. On the test set, all metrics remained equally strong, with AUC-ROC and AUC-PR at 100.00%, accuracy at 99.99%, precision at 99.95%, recall at 100.00%, and F1-score at 99.98%.




Figure 5: Performance Comparison of the HTGAD Fusion Model Across Training, Validation, and Test Datasets

4.6 Full Performance Comparison

Table 1 shows that XGBoost and HTGAD achieve near-perfect performance across training, validation, and test sets, with AUC-ROC, AUC-PR, accuracy, recall, and F1-scores all close to 100%. LSTM also performs strongly, achieving consistently high scores above 98%, though slightly lower than the top models. In contrast, GraphSAGE exhibits poor discriminative performance, with AUC values around 50% despite moderate accuracy, indicating limited effectiveness on the task.

Table 1: Full Classification Report Summary - All Models and All Splits

Model	AUC-ROC	AUC-PR	Accuracy	Precision	Recall	F1-Score
XGBoost Training	100.00%	99.99%	100.00%	99.98%	100.00%	99.99%
XGBoost Validation	100.00%	99.99%	99.99%	99.96%	100.00%	99.98%
XGBoost Test	100.00%	100.00%	99.99%	99.95%	100.00%	99.98%
GraphSAGE Training	50.95%	19.63%	80.63%	-	-	-
GraphSAGE Validation	49.63%	18.70%	81.03%	-	-	-
GraphSAGE Test	50.35%	19.53%	81.21%	-	-	-
LSTM Training	99.87%	99.49%	98.15%	92.79%	98.08%	95.36%
LSTM Validation	99.88%	99.51%	98.06%	92.16%	98.22%	95.09%
LSTM Test	99.89%	99.53%	98.41%	93.86%	98.07%	95.92%
HTGAD Training	100.00%	100.00%	100.00%	99.98%	100.00%	99.99%
HTGAD Validation	100.00%	99.99%	99.99%	99.96%	100.00%	99.98%
HTGAD Test	100.00%	100.00%	99.99%	99.95%	100.00%	99.98%

Confusion matrix analysis on the test set revealed stark differences in model behavior across 11,790 transactions. XGBoost correctly identified 9,574 normal and all 2,215 anomalous transactions with only one false positive. GraphSAGE predicted all transactions as normal, generating 2,215 false negatives. The LSTM produced 1,171 false positives and 1,921 false negatives, suggesting that five-transaction sequences provided some sequential context but not enough on their own. The HTGAD fusion model matched XGBoost, correctly identifying all anomalies with zero false negatives and only one false positive.

4.7 SHAP Feature Importance Analysis

SHAP analysis was applied to the XGBoost component to interpret which features most influenced anomaly risk predictions. Among all features, `transaction_type_purchase` had the highest mean absolute SHAP value at approximately 1.30, followed closely by `transaction_type_transfer` at about 1.28, indicating that the nature of the transaction is the most influential factor in identifying anomalous behavior. Temporal activity also played a significant role, with `hour_of_day` contributing approximately 0.55, suggesting that transactions occurring at unusual times increased the likelihood of being flagged as anomalous. Transaction amount contributed approximately 0.48, highlighting the importance of unusually high or irregular transaction values. Behavioral indicators such as `login_frequency` (0.30) and `transaction_type_scam` (0.16) further helped the model detect suspicious patterns. In contrast, geographic and demographic features, including `location_region_Europe`, `location_region_North America`, `ip_prefix`, and age groups, had minimal impact, with SHAP values close to 0.01 or lower.




Figure 6: SHAP Beeswarm Summary of Feature Contributions to Anomaly Risk Prediction.




Figure 7: Global Feature Importance for Anomaly Detection Based on Mean Absolute SHAP Values

5. DISCUSSION

The HTGAD framework delivered exceptionally high performance in detecting anomalies within Metaverse financial transactions. XGBoost consistently provided the strongest predictive power, correctly classifying nearly all normal and anomalous transactions, demonstrating the discriminative strength of transaction-level features including amount, type, session duration, and login frequency. LSTM effectively captured temporal dependencies, revealing sequential behavioral patterns that contributed significantly to anomaly identification. GraphSAGE, despite limited standalone performance due to the minimal node features employed (degree only), contributed critical structural information, including recognition of relational patterns such as coordinated wallet clusters and repeated transactional interactions that were not detectable by XGBoost or LSTM alone. The HTGAD fusion

successfully combined these complementary signals, appropriately weighting the stronger inputs while leveraging the unique contributions of all three models. The confusion matrix results are particularly informative: the HTGAD model achieved zero false negatives on the test set, which is the most critical outcome in financial fraud detection. False negatives represent undetected fraudulent transactions and carry the highest operational risk. That the hybrid framework preserved XGBoost's near-perfect anomaly recall while incorporating graph and temporal information validates the multi-perspective design. Table 2 contextualizes HTGAD against published approaches reported in the thesis literature review.

Table 2: Evaluation Metrics of Hybrid, Temporal, Graph, and Ensemble Models Across Financial and Synthetic Datasets

Model	Authors	AUC-ROC	F1-Score	Precision	Recall
AE-IF Hybrid	El Emary et al., 2026	0.952	0.870	0.880	0.860
IF + AE + XGBoost	Prakash et al., 2026	0.912	0.885	0.935	0.841
MDST-GNN	Chen et al., 2025	0.968	-	0.950	-
GeneralDyG	Yang et al., 2024	+3.6%	0.852	-	-
TGN-Mean	Kim et al., 2024	0.7747	-	-	-
Random Forest	Jumani & Raza, 2025	-	0.983	0.997	-
<i>HTGAD Fusion (Proposed)</i>	<i>Present Study</i>	<i>1.000</i>	<i>0.9998</i>	<i>0.9996</i>	<i>1.000</i>

Against this backdrop, the proposed HTGAD fusion model stands out, achieving an AUC-ROC of 1.000, an F1-score of 0.9998, a precision of 0.9996, and a recall of 1.000, indicating nearly flawless detection of anomalous instances across all evaluation criteria. Previous hybrid approaches, such as the AE-IF model by El Emary et al. (2026), reported an AUC-ROC of 0.952, and the IF plus AE plus XGBoost ensemble by Prakash et al. (2026) achieved an AUC-ROC of 0.912 with a recall of 0.841, suggesting that some anomalies were missed. Graph neural network approaches, including MDST-GNN (Chen et al., 2025), achieved an AUC-ROC of 0.968, while temporal models such as TGN-Mean (Kim et al., 2024) showed more modest performance at 0.7747. This comparison confirms the advantage of hybrid and graph-informed architectures in capturing complex relational patterns that individual models or simpler ensembles struggle to detect.

For financial service providers, cryptocurrency exchanges, and Metaverse-based marketplaces, the HTGAD framework provides actionable insights for enhancing fraud detection and risk management. Implementing the framework can help reduce operational risks, minimize false positives, and support data-driven decision-making, ultimately contributing to more secure and resilient digital financial operations. SHAP-based analysis further suggests that monitoring

efforts should prioritize high-impact variables, including transaction type, transaction amount, hour of day, and login frequency, thereby enabling more efficient allocation of investigative resources and risk mitigation strategies. Several limitations affected this research. The large dataset and complex features made training graph-based and temporal models computationally intensive. Missing or noisy data occasionally reduced model performance, and sparse or irregular sequences limited the performance of temporal and graph models. Class imbalance further complicated anomaly detection. The GraphSAGE component was limited to degree-only node features; incorporating richer node-level features, such as aggregated transaction volumes and behavioral metrics, would likely improve its predictive power and further strengthen the ensemble. The study also relied solely on a secondary simulated dataset from Kaggle, and generalizability to real-world Metaverse transaction data requires further investigation.

6. CONCLUSION

This study proposed and evaluated the Hybrid Temporal-Graph Anomaly Detection (HTGAD) framework for detecting anomalous and potentially fraudulent activities in Metaverse financial transactions. The framework combined XGBoost, LSTM, and GraphSAGE through a logistic regression meta-layer, delivering AUC-ROC values of 100%, F1-scores above 99.98%, and precision and recall exceeding 99.95% for anomalous transactions across training, validation, and test datasets. These results confirmed that integrating feature-based, temporal, and relational perspectives produced a highly reliable and balanced anomaly detection system that consistently outperformed individual models. Feature importance analysis using SHAP confirmed that transaction characteristics, timing, and behavioral metrics, including transaction type, transaction amount, hour of day, and login frequency, drove the model's predictions, while geographic and demographic features had minimal influence. This emphasized that transaction behavior and temporal patterns are the primary determinants of anomaly detection in Metaverse environments. Future work should enhance the GraphSAGE component by incorporating richer node-level features, such as aggregated transaction volumes and advanced network connectivity measures. Deploying HTGAD in operational environments for real-time anomaly detection could enhance the identification of high-risk and coordinated fraudulent activity. Incorporating advanced architectures, such as graph transformers or attention-based temporal models, could further refine the ensemble's ability to capture evolving and subtle patterns in dynamic Metaverse financial networks. Extending the hybrid methodology to other domains involving sequential and relational data, such as digital identity verification, supply chain monitoring, and decentralized marketplaces, would provide further opportunities to test the framework's generalizability and robustness.

REFERENCES

- Achuthan, K., Ramanathan, S., & Raman, R. (2025). Securing the metaverse: Machine learning-based perspectives on risk, trust, and governance. *International Journal of Information Management Data Insights*, 5(2), 100356. <https://doi.org/10.1016/j.jjime.2025.100356>
- Airlangga, G. (2024). Deep learning for anomaly detection and fraud analysis in blockchain transactions of the open metaverse. *Jurnal Informatika Ekonomi Bisnis*, 6(2). <https://doi.org/10.37034/inf.v6i2.865>
- Aysan, A. F., Gozgor, G., & Nanaeva, Z. (2024). Technological perspectives of metaverse for financial service providers. *Technological Forecasting and Social Change*, 202, 123323.

- Chen, S., Liu, Y., Zhang, Q., Shao, Z., & Wang, Z. (2025). Multi-distance spatial-temporal graph neural network for anomaly detection in blockchain transactions. *Advanced Intelligent Systems*, 7(8), 2400898.
- Cholevas, C., Angeli, E., Sereti, Z., Mavrikos, E., & Tsekouras, G. E. (2024). Anomaly detection in blockchain networks using unsupervised learning: A survey. *Algorithms*, 17(5), 201.
- El Emary, I. M. M. (2026). 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*, 3(1), 46–63.
- Ferretti, S., D'Angelo, G., & Ghini, V. (2025). Enhancing anti-money laundering frameworks: An application of graph neural networks in cryptocurrency transaction classification. *IEEE Access*, 13, 50201–50215.
- Huynh-The, T., et al. (2023). Blockchain for the metaverse: A review. *Future Generation Computer Systems*, 143, 401–419.
- Iftekhhar Rasul, S. M., et al. (2024). Detecting financial fraud in real-time transactions using graph neural networks and anomaly detection. *Journal of Economics, Finance and Accounting Studies*, 6(1), 131–142.
- Jumani, F., & Raza, M. (2025). Machine learning for anomaly detection in blockchain: A critical analysis, empirical validation, and future outlook. *Computers*, 14(7), 247.
- Kim, Y., Lee, Y., Choe, M., Oh, S., & Lee, Y. (2024). Temporal graph networks for graph anomaly detection in financial networks. *arXiv*. <https://doi.org/10.48550/ARXIV.2404.00060>
- Li, J., Lu, G., Gao, Y., & Gao, F. (2023). A smart contract vulnerability detection method based on multimodal feature fusion and deep learning. *Mathematics*, 11(23), 4823.
- Li, M. (2024). Meta-universe financial transaction anomaly detection and risk prediction based on machine learning. In *Proceedings of the 2024 2nd International Conference on Image, Algorithms and Artificial Intelligence (ICIAAI 2024)* (Vol. 115, pp. 117–129).
- Liu, T., et al. (2024). The role of transformer models in advancing blockchain technology: A systematic survey. *arXiv*. <https://doi.org/10.48550/ARXIV.2409.02139>
- Macak, M., et al. (2022). Process mining usage in cybersecurity and software reliability analysis: A systematic literature review. *Array*, 13, 100120.
- Mounnan, O., et al. (2024). A review on deep anomaly detection in blockchain. *Blockchain: Research and Applications*, 5(4), 100227.
- Pasaribu, J., Yudistira, N., & Mahmudy, W. F. (2024). Tabular data classification and regression: XGBoost or deep learning with retrieval-augmented generation. *IEEE Access*, 1–1.
- Pocher, N., et al. (2023). Detecting anomalous cryptocurrency transactions: An AML/CFT application of machine learning-based forensics. *Electronic Markets*, 33(1), 37.
- Prakash, S. (2026). Hybrid ensemble learning for anomaly detection in metaverse transactions using isolation forest, autoencoder, and XGBoost. *International Journal Research on Metaverse*, 3(1), 64–80.

- Ranjbaran, G., et al. (2025). C-SHAP: A hybrid method for fast and efficient interpretability. *Applied Sciences*, 15(2), 672.
- Sebt, M. V., Ghasemi, S. H., & Mehrkian, S. S. (2021). Predicting the number of customer transactions using stacked LSTM recurrent neural networks. *Social Network Analysis and Mining*, 11(1), 86.
- Siddamsetti, S. (2024). Anomaly detection in blockchain using machine learning. *Journal of Electrical Systems*, 20(3), 619–634.
- Umer, Q., et al. (2023). Ensemble deep learning-based prediction of fraudulent cryptocurrency transactions. *IEEE Access*, 11, 95213–95224.
- Wang, X., Guo, J., Luo, X., & Yu, H. (2024). DyHDGE: Dynamic heterogeneous transaction graph embedding for safety-centric fraud detection in financial scenarios. *Journal of Safety Science and Resilience*, 5(4), 486–497.
- Xu, B., et al. (2023). Study on the prediction of the uniaxial compressive strength of rock based on the SSA-XGBoost model. *Sustainability*, 15(6), 5201.
- Yang, X., Zhao, X., & Shen, Z. (2024). A generalizable anomaly detection method in dynamic graphs. *arXiv*. <https://doi.org/10.48550/ARXIV.2412.16447>
- Zhou, L., Zeng, Q., & Li, B. (2022). Hybrid anomaly detection via multihead dynamic graph attention networks for multivariate time series. *IEEE Access*, 10, 40967–40978.