

---

# Probabilistic Temporal Sampling for Anomaly Detection in Ethereum Networks

---

Anonymous Author(s)

Affiliation

Address

email

## Abstract

1 The rapid growth of the Ethereum network necessitates advanced anomaly detec-  
2 tion techniques to enhance security, transparency, and resilience against evolving  
3 malicious activities. While there have been significant strides in anomaly detection,  
4 they often fall short in capturing the intricate spatial-temporal patterns inherent  
5 in blockchain transactional data. This study presents a scalable framework that  
6 integrates Graph Convolutional Networks (GCNs) with Temporal Random Walks  
7 (TRW) specifically designed to adapt to the complexities and temporal dynamics  
8 of the Ethereum transaction network. Unlike traditional methods that focus on  
9 detecting specific attack types, such as front-running or flash loan exploits, our  
10 approach targets time-sensitive anomalies more broadly—detecting irregularities  
11 such as rapid transaction bursts, anomalous token swaps, and sudden volume spikes.  
12 This broader focus reduces reliance on pre-defined attack categories, making the  
13 method more adaptable to emerging and evolving malicious strategies. To ground  
14 our contributions, we establish three theoretical results: (1) the effectiveness of  
15 TRW in enhancing GCN-based anomaly detection by capturing temporal dependen-  
16 cies, (2) the identification of weight cancellation conditions in the anomaly  
17 detection process, and (3) the scalability and efficiency improvements of GCNs  
18 achieved through probabilistic sampling. Empirical evaluations demonstrate that  
19 the TRW-GCN framework outperforms state-of-the-art Temporal Graph Attention  
20 Networks (TGAT) in detecting time-sensitive anomalies. Furthermore, as part  
21 of our ablation study, we evaluated various anomaly detection techniques on the  
22 TRW-GCN embeddings and found that our proposed scoring classifier consistently  
23 achieves higher accuracy and precision compared to baseline methods such as  
24 Isolation Forest, One-Class SVM, and DBSCAN, thereby validating the robustness  
25 and adaptability of our framework.

## 26 1 Introduction

27 The Ethereum network is a dynamic and complex ecosystem, characterized by high-frequency  
28 transactions, time-sensitive interactions, and evolving patterns of fraudulent activity. Anomalous  
29 behaviors such as flash loans, front-running attacks, and MEV (Miner Extractable Value) bots pose  
30 significant threats to the security and integrity of the network. These behaviors often unfold over  
31 time, making it essential to account for temporal correlations in transaction patterns for effective  
32 anomaly detection. The dynamic nature of Ethereum presents unique challenges that cannot be fully  
33 addressed using static graph analysis or traditional machine learning approaches.

34 Graph Convolutional Networks (GCNs) have emerged as a transformative tool in the domain of graph-  
35 structured data representation. Their ability to encapsulate both local and global graph structures has  
36 paved the way for their application in diverse fields. While traditional GCNs have shown remarkable  
37 potential in handling static graph structures, their application to dynamic graphs introduces new

38 challenges and opportunities. In order to extend GCNs to dynamic graphs, it is crucial to understand  
 39 how learning on dynamic graphs works, which is a relatively recent area of research. There have  
 40 been studies which investigate discrete-time graphs represented as a sequence of graph snapshots  
 41 (1) (2) (3). Also several continuous-time approaches have been presented (4) (5) (6) (7) (8), where  
 42 continuous dynamic graphs means that edges can appear at any time (8) (9).

43 Also, the topic of anomaly detection in Blockchain has received considerable attention. For example,  
 44 in Ethereum, the unexpected appearance of particular subgraphs has implied new malware (10).  
 45 Anomaly detection in blockchain transaction networks is an emerging area of research in the cryp-  
 46 tocurrency community (11). Wu et al. (12) investigated phishing detection in blockchain network  
 47 using unsupervised learning algorithms. Ofori-Boateng et al. (13) have also discussed topological  
 48 anomaly detection in multilayer blockchain networks. Given that the Ethereum network witnesses  
 49 dynamically evolving transaction patterns, it becomes imperative to account for the temporal se-  
 50 quences and correlations of transactions. Unlike general-purpose graph neural networks, TRW-GCN  
 51 is a domain-specific framework tailored to the Ethereum network’s unique dynamics, see Table 1 for  
 52 comparison. By leveraging temporal features and dynamic embeddings, our approach enables the  
 53 detection of time-sensitive anomalies such as flash loans and MEV bots, with minimal computational  
 54 complexity. Our research offers several contributions:

55 **Enhanced Anomaly Detection Effectiveness:** Our model leverages TRW in tandem with GCN  
 56 to improve anomaly detection effectiveness. This integration improves the detection of anomalies  
 57 in the Ethereum transaction network by effectively leveraging temporal information embedded  
 58 within transaction patterns. The model’s ability to analyze temporal correlations allows it to identify  
 59 anomalies that traditional methods often overlook.

60 **Efficiency in Sampling Representative Nodes:** Given the substantial size and continuous growth  
 61 of the Ethereum blockchain, efficient sampling methods are essential. Our TRW-GCN provides  
 62 a solution that balances accuracy with computational efficiency. Many temporal graph learning  
 63 frameworks face performance bottlenecks when applied to densely connected graphs; for instance,  
 64 models such as TGAT (4) and AddGraph (14) incorporate temporal dynamics but often come with  
 65 high computational costs and are sensitive to the quality of temporal features, which can limit their  
 66 applicability to Ethereum’s specific requirements, whereas TRW-GCN prioritizes edges based on  
 67 their timestamps, enabling the model to capture time-sensitive relationships without the overhead of  
 68 attention mechanisms used in models like TGAT. See Table 1 for comparison.

69 **Detecting Patterns Leading to Sophisticated Attacks:** While existing works like "Flash Boys 2.0"  
 70 (15) and "Combating Front-Running in Smart Contracts" (16) which focus on detecting front-running  
 71 attacks specifically, our approach targets time-sensitive anomalies more broadly. These anomalies  
 72 include behaviors that may precede or indirectly relate to specific exploits, such as Front-Running  
 73 Transactions, Flash Loan Exploits, High-Frequency Token Swaps, and Irregular Contract Interactions,  
 74 see Table 5 for definitions. By identifying these timing-dependent irregularities, our work addresses a  
 75 wider range of anomalous behaviors that are indicative of potential security threats.

Table 1: Comparison of TRW-GCN with Existing Temporal GNNs in Blockchain context

Aspect	TGAT	AddGraph	TRW-GCN
<b>Temporal Modeling Mechanism</b>	Temporal attention on time-encoded node embeddings	Time-decay functions over temporal edges	Temporal random walks to construct time-aware neighborhoods
<b>Domain Specialization (Ethereum)</b>	General-purpose model with time-aware positional encodings and attention mechanisms	General-purpose, may underperform TGAT, Less interpretable than attention models	Tailored for Ethereum with attention to domain-specific phenomena (e.g., flash loans, MEV, wash trading)
<b>Anomaly Type Detection</b>	Detects broad irregularities, limited granularity, Heavy computation due to multi-head	Captures gradual shifts, not sharp transaction bursts	Detects fine-grained, time-sensitive anomalies like front-running and high-frequency exploits
<b>Robustness to Transaction Bursts</b>	Limited; signal may be diluted by attention weights	Time-decay may smooth over bursts	High; TRW preserves burst patterns in short temporal windows
<b>Real-World Applicability to Ethereum</b>	Rare in blockchain studies; lacks deployment cases	Not used in Ethereum networks	Demonstrates superior results in transaction-based anomaly detection

## 76 2 Model Design

77 GCNs are a pivotal neural network architecture crafted specifically for graph-structured data. Through  
 78 the use of graph convolutional layers, we seamlessly aggregate information from neighboring nodes  
 79 and edges to refine node embeddings. In enhancing this mechanism, we incorporate probabilistic

80 sampling, which proves particularly adept in analyzing the vast Ethereum network. The incorporation  
 81 of TRW adds a rich layer to this framework. TRW captures the temporal sequences in Ethereum  
 82 transactions and not only focuses on nodes’ spatial prominence but also considers the transactional  
 83 chronology. Here, ’time’ is conceptualized based on the sequence and timestamps of Ethereum  
 84 transactions, leading to a dynamically evolving, time-sensitive representation of the network.

85 Here, graph is represented as  $G = (V, E)$ , where  $V$  is the set of nodes (vertices) and  $E$  is the set  
 86 of edges connecting the nodes. Each node  $v_i$  in the graph is associated with a feature vector  $F_i$ ,  
 87 and  $F \in \mathbb{R}^{|V| \times 4}$  represents a feature matrix of size 4. Aggregation is a process to combine the  
 88 feature vectors of neighboring nodes using an adjacency matrix  $A$  to capture graph connectivity. To  
 89 enable information propagation across multiple layers, the graph convolution operation is performed  
 90 iteratively through multiple graph convolutional layers (GCLs). The output of one layer serves as  
 91 the input to the next layer, allowing the propagation of information through the network. The node  
 92 representations are updated layer by layer, allowing information from neighbors and their neighbors  
 93 to be incorporated into the node features. The parameters  $W^l$  are learned during the training process  
 94 to optimize the model’s performance on a specific graph-based task. GCNs often consist of multiple  
 95 layers, where each layer iteratively updates the node representations:

$$h_i^{(l)} = \text{Activation} \left( W^{(l)} \text{Aggregate} \left( h_j^{(l-1)} \mid j \in N(i) \right) \right) \quad (1)$$

96 Here,  $h_i^{(l)}$  is the representation of node  $i$  at layer  $l$ , and  $h_j^{(l-1)}$  is the representation of neighboring  
 97 node  $j$  at the previous layer ( $l-1$ ). The final layer is usually followed by a global pooling operation to  
 98 obtain the graph-level representation. The pooled representation is then used to make predictions.  
 99

## 100 2.1 Incorporating TRW into GCN

101 The TRW-enhanced GCN creates a multidimensional representation that captures both the structural  
 102 intricacies and time-evolving patterns of transactions. Such an approach requires meticulous math-  
 103 ematical modeling to substantiate its efficacy, and exploring the depths of this amalgamation can  
 104 reveal further insights into the temporal rhythms of the Ethereum network.

### 105 Temporal Random Walk (TRW)

106 Given a node  $i$ , the probability  $P_{ij}$  of moving to a neighboring node  $j$  can be represented as:

$$P_{ij} = \frac{\omega_{ij}}{\sum_k \omega_{ik}} \quad (2)$$

107 where  $\omega_{ij}$  is the weight of the edge between node  $i$  and  $j$ , and the denominator is the sum of weights  
 108 of all edges from node  $i$ . In a TRW, transition probabilities take into account temporal factors. Let’s  
 109 define the temporal transition matrix  $T$  where each entry  $T_{ij}$  indicates the transition probability from  
 110 node  $i$  to node  $j$  based on temporal factors.

$$T_{ij} = \alpha \times A_{ij} + (1 - \alpha) \times f(t_{ij}) \quad (3)$$

111 where  $A_{ij}$  is the original adjacency matrix’s entry for nodes  $i$  and  $j$ .  $\alpha$  is a weighting parameter.  $f_{ij}$   
 112 is a function of the temporal difference between node  $i$  and node  $j$ . The temporal weighting function  
 113 could be defined as an exponentially decaying function:

$$f(t_{ij}) = \exp(-\gamma \cdot t_{ij}) \quad (4)$$

114 where  $\gamma > 0$  is a decay hyperparameter that controls how sensitive the model is to temporal  
 115 differences.  $f(t_{ij}) \in [0, 1]$ , with values closer to 1 for temporally close nodes and closer to 0 for  
 116 nodes that are far apart in time. Given this temporal transition matrix  $T$ , a normalized form  $\tilde{T}$  can be  
 117 used for a GCN layer:

$$\tilde{T} = \tilde{D}_T^{-1} T \quad (5)$$

118 where  $\tilde{D}_T$  is the diagonal degree matrix of  $T$ . To incorporate the TRW’s temporal information into  
 119 the GCN, we can modify the original GCN operation using  $\tilde{T}$ :

$$h^{(l+1)} = \sigma \left( \tilde{D}_T^{-\frac{1}{2}} \tilde{T} \tilde{D}_T^{-\frac{1}{2}} h^{(l)} W^{(l)} \right) \quad (6)$$

## 120 2.2 Effect on Anomaly Detection

121 The embeddings from a GCN (post TRW influence) should be more sensitive to recent behaviors  
 122 and patterns. When these embeddings are passed to a classifier, clustering and scoring algorithms  
 123 like DBSCAN, OCSVM, ISOLATION FOREST, and LOF, anomalies that are based on recent or  
 124 time-sensitive behaviors are more likely to stand out. In our work, the term "anomaly" refers to  
 125 patterns that are statistically uncommon or divergent from the norm based on the features learned  
 126 by our model. These uncommon patterns, while not definitively erroneous, are of interest because  
 127 they deviate from typical behavior. In the context of Ethereum transactions, such deviations indicate  
 128 suspicious activities, novel transaction patterns, or transaction bursts.

129 While we here provide insight and mathematical proofs, the true value of TRW in improving GCN  
 130 over traditional sampling is empirical. We will compare the performance of GCN with and without  
 131 TRW on a temporal dataset to see tangible benefits (see appendix B.5). Here is how temporal weights  
 132 are applied:

- 133 1. Node Features are weighted by time: When updating the node features through the matrix  
 134 multiplication, nodes that are temporally closer influence each other more, allowing recent  
 135 patterns to be highlighted.
- 136 2. Temporal Relationships are captured: The modified node features inherently capture tempo-  
 137 ral relationships because they aggregate features from temporally relevant neighbors.
- 138 3. Higher Sensitivity to recent anomalies: With temporal weighting, anomalies that have  
 139 occurred recently will be more pronounced in the node feature space.

140 **Theorem 1:** Let  $G = (V, E)$  be a graph with node features  $h_i \in \mathbb{R}^d$  for  $i \in V$ , and let a GCN  
 141 generate embeddings through neighbor aggregation. Incorporating TRW, represented by a temporal  
 142 weight matrix  $T$ , into the aggregation mechanism enhances the effectiveness of detecting temporally  
 143 influenced anomalies. Specifically, if  $T$  encodes temporal transitions such that  $T_{ij} \neq 1$  for all  $i, j$ ,  
 144 the feature representation  $h_i^{(l+1)}$  for an anomalous node  $n$  differs significantly from the non-temporal  
 145 case:

$$\|h'_n - h_n\|_2 > \delta, \quad (7)$$

146 for some sensitivity threshold  $\delta > 0$ , where  $h_n$  is the embedding without TRW and  $h'_n$  is the  
 147 embedding with TRW.

### 148 Proof.

149 Anomaly detection is the task of distinguishing outliers from normal data points in a given feature  
 150 space. If we have an anomaly score function  $s : \mathbb{R}^d \rightarrow \mathbb{R}$ , we can detect anomalies by:  $s(v) > \delta$   
 151 Where  $\delta$  is a threshold, and  $\underline{v}$  is a vector in the feature space.

152 A GCN produces node embeddings (or features) by aggregating information from a node's neighbors  
 153 in the graph. Let's express this aggregation for a single node using a simple form of a GCN layer:

$$h_i^{(l+1)} = \sigma \left( \sum_{j \in \text{Neighbors}(i)} W h_j^{(l)} \right) \quad (8)$$

154 Where  $\underline{h}_i^{(l)}$  is the feature of node  $\underline{i}$  at layer  $\underline{l}$ , and  $\underline{W}$  is the weight matrix.

156 **Incorporating TRW:** With a temporal random walk, the aggregation process is influenced by time,  
 157 so the aggregation becomes:

$$h_i^{(l+1)} = \sigma \left( \sum_{j \in \text{Neighbors}(i)} T_{ij} W h_j^{(l)} \right) \quad (9)$$

158 Where  $T_{ij}$  is the temporal transition probability from node  $\underline{j}$  to node  $\underline{i}$ . Let's assume a node with an  
 159 anomaly will have a different feature vector from the nodes without anomalies. For simplicity, let's  
 160 use the Euclidean distance as the anomaly score:  $s(v) = \|v - \mu\|$  where  $\mu$  is the mean vector of all  
 161 node features. Given a temporal anomaly (an anomaly that's influenced by recent events), using TRW  
 162 will result in a modified feature vector for the anomalous node. Let's consider two scenarios:  
 163

- 164 1. GCN without TRW: For an anomalous node  $\underline{n}$ , its feature vector is:  $h_n = \sigma \left( \sum_j W h_j \right)$

165 2. GCN with TRW: For the same anomalous node  $\underline{n}$ , it becomes:  $h'_n = \sigma \left( \sum_j T_{nj} W h_j \right)$

166 If the anomaly is temporally influenced, then  $h'_n$  should be significantly different from  $h_n$  due  
167 to the weights introduced by  $T_{ni}$  (see Theorem 2 for weight cancellation). In the context of our  
168 anomaly score function:  $s(h'_n) - s(h_n) > \delta$  where  $\delta$  is a value indicating the sensitivity of the  
169 temporal context; we will use this later in our scoring method. If the anomaly is truly temporally  
170 influenced, this difference will be significant, and thus, the GCN with TRW will have a higher  
171 likelihood of detecting the anomaly. From the linear algebra perspective, the effect of TRW on a  
172 GCN for anomaly detection is evident in how node features are aggregated. The temporal weights  
173 (from  $T_{ij}$ ) make the GCN more sensitive to temporal influences, making it more adept at detecting  
174 anomalies. The theoretical result in Theorem 1 holds for any d-dimensional feature vector, including  
175 the 10-dimensional vectors used in the empirical section.  
176

177 **Theorem 2:** Let  $\mathbb{R}^m$  be a vector space, and let  $h_n \in \mathbb{R}^m$  represent a feature vector. Define a  
178 temporal transformation matrix  $T_{nj} \in \mathbb{R}^{m \times m}$ , where each entry  $t_{ij}$  encodes the temporal weights.  
179 Let  $h'_n = T_{nj} h_n$  be the transformed feature vector.

180 If the transformation matrix  $T_{nj}$  exhibits symmetric or complementary weight patterns that cause  
181 significant weight cancellation, the difference between the transformed and original vectors,  $\|h'_n -$   
182  $h_n\|_2$ , will be insufficient to surpass a given anomaly detection threshold  $\delta > 0$ .  
183 This proof is given in appendix A.

184 **Theorem 3:** Let  $G = (V, E)$  represent an Ethereum transaction graph with  $|V| = N$  nodes  
185 (accounts) and  $|E| = M$  edges (transactions). Let  $X \in \mathbb{R}^{N \times d}$  denote the feature matrix for the  
186 nodes,  $A \in \mathbb{R}^{N \times N}$  the adjacency matrix representing transaction relationships, and  $Y \in \{0, 1\}^N$  the  
187 binary labels indicating specific account behaviors. Probabilistic random walk sampling, defined by a  
188 sampling matrix  $P$ , improves the performance of a GCN for the task of predicting node labels  $Y$  in  
189 the context of Ethereum networks.  
190 This proof is given in appendix B.

## 191 3 Empirical Analysis

192 In this section, not only we provide details about the empirical analysis and evaluation methods, but  
193 also provide supporting information for readers to follow the experiments.

### 194 3.1 Datasets, Materials and Methods

195 We provide datasets and the code in the github link [https://github.com/stefankam/temporal-spatial-](https://github.com/stefankam/temporal-spatial-anomaly-detection)  
196 [anomaly-detection](https://github.com/stefankam/temporal-spatial-anomaly-detection). We run the code on our department server running Linux equipped with a single  
197 GPU (NVIDIA A100 80GB PCIe), and 251Gi RAM.

198 Creating a complete transaction graph for all Ethereum blocks would be a computationally intensive  
199 task, as it would involve processing and storing a large amount of data. However, in the supplemental  
200 material we provide the code to generate a transaction history graph for a range of 100-1000 blocks.  
201 We further incorporate spatial and temporal node features to capture temporal aspects more explicitly:

202 **incoming\_value\_variance:** Variance of the transaction values received by the node. This metric  
203 quantifies the spread or dispersion of incoming transaction amounts, providing insight into the  
204 consistency or variability of funds received. **outgoing\_value\_variance:** Variance of the transaction  
205 values sent by the node. **activity\_rate:** The activity rate of a node represents the total number of  
206 transactions (both incoming and outgoing) divided by the duration (in terms of blocks). It indicates  
207 the frequency of interactions involving the node over a specific period. **change\_in\_activity:** The  
208 change in activity refers to the difference in the number of transactions of the current block compared  
209 to the previous block for a given node. This metric captures fluctuations or deviations in transaction  
210 behavior over consecutive blocks. **time\_since\_last:** Time since the last transaction involving the node,  
211 measured as the difference between the current block number and the block number of the node's  
212 most recent transaction. It provides insights into the recency of activity associated with the node.  
213 **tx\_volume:** Total transaction volume associated with the node, calculated as the sum of incoming and  
214 outgoing transaction values. This metric represents the overall magnitude of financial transactions  
215 involving the node. **frequent\_large\_transfers:** Indicator variable identifying addresses engaged

216 in frequent and large transfers. Nodes meeting specific thresholds for both transaction frequency  
217 and volume are flagged. **gas\_price**: Additional feature relevant for MEV detection, representing the  
218 gas price paid for transactions. Gas price fluctuations can signal potential MEV activities such as  
219 frontrunning or transaction ordering strategies. **token\_swaps**: Another feature for MEV detection,  
220 indicating involvement in token swaps or trades on decentralized exchanges (DEXs). Analysis  
221 of token swap transactions can reveal arbitrage opportunities or manipulative behavior by MEV  
222 bots. **smart\_contract\_interactions**: Feature identifying transactions interacting with known DeFi  
223 protocols or smart contracts. MEV bots may exploit vulnerabilities or manipulate protocol behaviors.

### 224 3.2 TRW-GCN combined method to detect anomalies

225 To apply graph convolutional layers to the blockchain data for aggregating information from neigh-  
226 boring nodes and edges, we'll use the PyTorch Geometric library. This library is specifically designed  
227 for graph-based data and includes various graph neural network layers, including graph convolutional  
228 layers. Note that training and testing a graph neural network on Ethereum dataset would require  
229 significant computational resources, as currently, the Ethereum network possesses about 20 million  
230 blocks, which are connected over the Ethereum network. In this study, we provide the transaction  
231 history within a specified range of 1000 blocks; we believe, adding blocks do not add any advantage.

232 In Algorithm 1, we intend to compare the anomaly detection of full- and sub-graphs (sampling  
233 using TRW). The graph convolution operation combines the features of neighboring nodes to update  
234 the representation of a given node. As node features, we input the 10 features indicated in 3.1 as  
235 vector representation; considering 20 hidden layers, 100 epochs,  $lr=0.01$ ,  $num\_walks=10$ , and  
236  $walk\_length=100$ , the resulting output vector aggregates information from all neighboring nodes.  
237 By using the nodes from TRW for training, the GCN will be more attuned to the time-dependent  
238 behaviors, leading to better detection of sudden spikes in transaction volume or unusual contract  
239 interactions that occur in quick succession. In our experiments, we employ TRW to sample nodes  
240 from the entire graph, ensuring that the graph's integrity is maintained. Here's how it can be done:

- 241 1. **Perform TRWs to Sample Nodes for Training:** The TRWs provide sequences of nodes  
242 representing paths through the Ethereum network graph. Nodes appearing frequently in  
243 these walks are often involved in recent temporal interactions.
- 244 2. **Train the GCN with the Sampled Nodes:** Instead of using the entire Ethereum network  
245 graph for training, use nodes sampled from the TRWs. This approach tailors the GCN to  
246 recognize patterns from the most temporally active parts of the Ethereum network.

247 Using the GCN with TRW combined method, one can achieve 1) Anomalies detected, 2) Training  
248 efficiency, and 3) Quality of embedding. The integration of TRW with GCNs offers a novel approach  
249 for generating embedding that capture both spatial and temporal patterns within the Ethereum network.  
250 These embedding are vital for understanding the underlying transaction dynamics and for effectively  
251 detecting anomalous activities. To evaluate the potential of the TRW-GCN methodology, we employ  
252 four distinct machine learning techniques: DBSCAN, SVM, Isolation Forest (IsoForest), and Local  
253 Outlier Factor (LOF). Wu et al. (12) indicated that they have obtained more than 500 million  
254 Ethereum addresses and 3.8 billion transaction records. However, only 1259 addresses are labeled as  
255 phishing addresses collected from EtherScamDB, which implies an extreme data imbalance as the  
256 biggest obstacle for phishing detection, therefore they used unsupervised learning detection method.  
257 We similarly use unsupervised learning for detection in our TRW-GCN algorithm.

258 The extensive use of these four diverse methods allows us to validate the efficacy of the TRW-GCN  
259 framework. The high anomaly detection rates in Figure 1 by clustering methods underscores the  
260 importance of algorithm selection. As easily observed, using the embedding generated by TRW-GCN  
261 in SVM method significantly improves anomaly detection, however other methods do not show any  
262 improvement in anomaly detection (averaged over 10 runs); the enhanced detection capabilities in  
263 SVM could be attributed to the TRW's ability to encapsulate temporal sequences and correlations of  
264 transactions. In Table 2, we compare these methods in terms of their precision, recall and F-score and  
265 compare with the outcome of SVM and IsoForest methods implemented in (12) (note that this paper  
266 focuses on Phishing detection in Ethereum Network, and is different from our dynamic approach in  
267 temporal anomaly detection). Our models are marking many data points as anomalies; precision stays  
268 relatively high, but low F-score. Algorithms like DBSCAN, LOF, Isolation Forest are unsupervised,  
269 so they often overpredict if not-cluster-fit noise is high or parameter tuning is off. DBSCAN is very  
270 sensitive to eps. Isolation Forest depends on contamination, and LOF is sensitive to  $n\_neighbors$ .

Algorithm 1: TRW- GCN combined method to detect anomalies

---

**Steps:**

1. Load and Preprocess the graph  $G$ .
2. **For** each walk  $k = 1$  to  $\text{num\_walks}$ :  
 $W = \{\text{w\_k for k in range}(1, \text{num\_walks}+1)\}$  for  $\text{w\_k}$  in  $\text{temporal\_random\_walk}(k)$   
// Aggregate walks in  $W$
- End**
3. **For** training step:  
 $F = \text{torch.stack}([\text{f}(v_i) \text{ for } v_i \text{ in } V], \text{dim}=0)$   
 $A = \text{nx.to\_numpy\_matrix}(\text{graph}, \text{nodelist}=V)$   
 $M_{TRW}, M = \text{GCN}(\text{in\_channels}, \text{hidden\_channels}, \text{out\_channels})$   
 $\text{train}(M_{TRW}, F, A)$  if use\\_TRW else  $\text{train}(M, F, A)$   
// Training using sampled-graphs
- End**
4. Apply DBSCAN, One-Class SVM, IsoForest, and LOF on embeddings from the trained GCN model  $M$  to obtain anomalies.

---

Algorithm 2: A Score-based anomaly detection associated with time-dependent behaviors

---

**Steps:**

1.  $G' = G(V, E)$  where  $E$  has node attributes.
2.  $X = [x_1, x_2, \dots, x_n]$  for  $n \in V$ .
3. GCNModel with layers:  
 $\text{in\_channels} \rightarrow \text{hidden\_channels} \rightarrow \text{out\_channels}$
4.  $\text{TRW}(G', \text{start}, \text{length})$  returns walk  $W$  and timestamps  $T$
5. **For** each walk  $i = 1$  to  $\text{num\_walks}$ :  
 $\text{All\_Walks} = \bigcup_{i=1}^{\text{num\_walks}} \text{TRW}(G', \text{random\_node}, \text{walk\_length})$   
// Node Sampling via TRW
- End**
6. Node Frequency Computation:  
 $\text{freq}(v) = \frac{\text{occurrences of } v \text{ in All\_Walks}}{\text{max occurrences in All\_Walk}}$  for  $v \in V$ .
7. Anomaly Score Computation:  
 $S(v) = \frac{(\text{emb}(v)_{\text{latest}} - \mu(\text{emb}(v)))}{\sigma(\text{emb}(v))} \times \text{freq}(v)$   
where  $\text{emb}$  is the node embedding,  $\mu$  is the mean, and  $\sigma$  is the standard deviation; anomalies are detected when  $S(v) > \text{threshold } \delta$ .

---

271 It is also interesting to find out which node features mainly contribute to anomaly detection; we  
272 show this in Figure 2. As illustrated by different colors, the feature 3-6 namely `activity_rate`,  
273 `change_in_activity`, `time_since_last` (mainly the temporal features) are the drivers of frequent anom-  
274 alies (with dark blue colors), while `tx_volume` and `frequent_large_transfers` (with green colors) also  
275 produce anomalies but less frequently. Although we have obtained good insights into the method  
276 effectiveness to detect time-dependent patterns and features, but we should look for more precise and  
277 less prone to error detection method.




Figure 1: A comparison (mean, std) of 4 detection models namely `dbscan`, `svm`, `isoforest` and `lof` between full-graph and sub-graph with TRW sampling. Using TRW-GCN clearly improves SVM in anomaly detection; other methods do not seem to be improved.

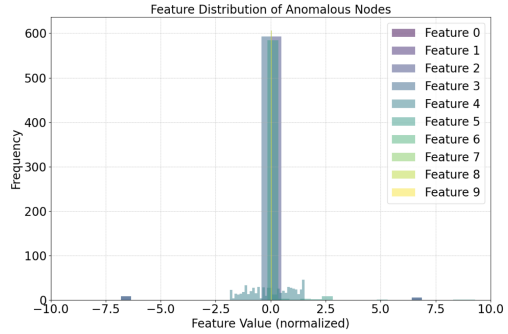


Figure 2: Feature distribution where Blue and Green colors: `activity_rate`, `change_in_activity`, and `times_since_last` show highest frequencies.

Table 2: Comparison of Precision/Recall/F-score of 4 methods with/out-TRW

Method	Prec.(w-T)	Rec.(w-T)	F-S.(w-T)	Prec.(o-T)	Rec.(o-T)	F-S.(o-T)	Prec.(12)	Rec.(12)	F-S.(12)
DBSCAN	0.799	0.485	0.604	0.799	0.485	0.604			
SVM	0.799	0.438	0.563	0.796	0.333	0.458	0.927	0.893	0.908
IsoForest	0.795	0.094	0.163	0.796	0.094	0.163	0.821	0.849	0.835
LOF	0.815	0.096	0.167	0.812	0.096	0.167			

### 278 3.3 Score-based anomaly pattern

279 While traditional methods compute anomaly scores based on the relative position or density of data  
280 points in the feature space, we need a method to be more focused on temporal dynamics, tracking  
281 the evolution of each node’s embedding over time and weighing it by the node’s frequency in the  
282 graph. To adapt the code to pick up anomalous patterns associated with time-dependent behaviors,  
283 the algorithm should be equipped to recognize such patterns. Hence, we augment the node features  
284 to capture recent activities with time features as explained in 3.1 dataset section, and after obtaining  
285 node embedding from the GCN, compute the anomaly score for each node based on its temporal  
286 behavior. The simplest way to achieve this is by computing the standard deviation of the node’s  
287 feature over time and checking if the latest data point deviates significantly from its mean. This was  
288 initially discussed in Theorem 1, with weight cancellation argument in Theorem 2.

289 We explain all the steps in Algorithm 2. Initially, we define the node features to capture recent  
 290 activities. After training the GCN and obtaining the embedding, we compute an anomaly score based  
 291 on how much the recent transaction volume (the latest day in our case) deviates from the mean. We  
 292 then use a visualization function to display nodes with an anomaly score beyond a certain threshold  
 293 (in this case, we've used a z-score threshold of 2.0 which represents roughly 95% confidence).




Figure 3: Anomaly detection in (left) 100 blocks with 6 features, (middle) 100 blocks with 10 features, (right) the anomalous addresses where the time-sensitive associated ones are hashed green.

294 In Figure 3, black points represent the vast majority of nodes in our specified Ethereum network  
 295 dataset; they signify regular non-anomalous Ethereum addresses. Cluster of points inside and  
 296 around the blue circle represent groupings of Ethereum addresses or contracts that have had frequent  
 297 interactions with each other. The density or proximity of points to each other indicates how closely  
 298 those addresses or contracts are related. Red points would represent the nodes that have been flagged  
 299 as anomalous based on their recent behavior. The code identifies them by computing an anomaly  
 300 score, and those exceeding a threshold are colored red. In the left graph, there are just 20 nodes  
 301 detected as anomaly in 100 blocks where we used 6 structural features in our detection algorithm,  
 302 while in the middle graph, we used 10 features to detect anomalies in the same 100 blocks, and  
 303 12 more suspicious addresses are detected, hashed in green in the right figure. This signifies the  
 304 importance of temporal feature selection, as by adding 4 temporal features we would be able to  
 305 detect missing anomalies. We checked these addresses in Ether explorer website <https://etherscan.io> ,  
 306 and found the corresponding labels such as MEV Bot, Metamask: Swap, Uniswap, Wrapped Ether,  
 307 Rollbit, Blur: Bidding, which are mainly time-sensitive transactions or contracts, see next section  
 308 for explanation on what is normal versus anomaly. In Table 3, we explain the types of such detected  
 309 anomalies and the associated addresses. This is a proof of cross-checking with the ground-truth .

Table 3: Some types of detected anomalies

Ethereum addresses for anomalies detected from Figure 3	Ground Truth : cross-check with <a href="https://etherscan.io/">https://etherscan.io/</a>
0x6F1cDbBb4d53d226CF4B917bF768B94acbAB6168	MEV Bot; certain activities may be considered harmful
0x3fC91A3afd70395Cd496C647d5a6CC9D4B2b7FAD	Uniswap (users to swap various ERC-20 tokens)
0x881D40237659C251811CEC9c364ef91dC08D300C	Metamask Swap router
0x0000000000A39bb272e79075ade125fd351887Ac	Flashloan; Detecting involves transactions with large token volumes

Table 4: TRW-GCN versus TGAT for eth\_latest\_100\_block file, and z-score threshold of 2.0.

Model	Accuracy / # Anomalies detected
TRW-GCN	94.5% / 20
TGAT	85.3% / 23

### 3.4 Normal versus Anomaly, Baseline algorithm, Algorithm complexity, and the Ground truth

312 In Ethereum, what may be considered normal or anomalous behavior can vary depending on various  
 313 factors such as market conditions, network activity, and the specific use cases of different addresses  
 314 or smart contracts. Time-sensitive irregularities in Ethereum transactions refer to anomalies that  
 315 occur within specific time frames or exhibit patterns that are indicative of immediate or rapid actions.  
 316 These irregularities may include instances of rapid buying or selling of assets, front-running other  
 317 traders, MEV activities, flash loan exploits, or token swaps executed within short time intervals.  
 318 Identifying these irregularities requires analyzing transactional data in real-time or within narrow  
 319 time windows to capture anomalous behaviors as they occur. See Table 5 for a list of time-sensitive  
 320 items in Ethereum network including transactions, contracts, and platform activities. Our objective  
 321 is to identify such instances; upon identifying suspicious transactions, our approach advocates for



322 further investigation. In Table 3, we cross-reference the transaction details with etherscan.io (which  
 323 represents a source for ground truth, where one finds more information about an anomaly).

Table 5: some time sensitive items on Ethereum network and their definitions

Time sensitive items	Definitions
MEV Bot	MEV refers to the additional profit that miners can extract from the Ethereum network by reordering, censoring, or including transactions in blocks. The timing of transactions and block mining can affect the potential profit extracted by MEV bots. MEV can affect fairness and efficiency of the Ethereum network.
Metamask: Swap Uniswap	Uniswap is a decentralized exchange (DEX) protocol on Ethereum, and swaps conducted through MetaMask can be time-sensitive, especially considering the volatility of cryptocurrency prices and liquidity on Uniswap.
Flashloan	Flash loans are uncollateralized loans that must be borrowed and repaid within a single transaction block. These loans are often used for arbitrage, liquidations, or other trading strategies that require rapid execution.
Wrapped Ether (WETH)	Wrapped Ether (WETH) is an Ethereum token pegged to the value of Ether (ETH). Transactions involving WETH can be time-sensitive, especially if they're related to trading, liquidity provision, or token swaps.
Token Launches and Airdrops	Token launches and airdrops often have predefined distribution schedules or timeframes during which users can claim or receive tokens.
Smart Contract Exploits	Exploiting vulnerabilities in smart contracts often requires precise timing to execute malicious transactions before vulnerabilities are patched or mitigated.

324 Similar to the papers by Wu et al. (12), Zhang et al. (16), and Feng et al. (17), as baseline algorithms  
 325 for comparison, common unsupervised methods such as Isolation Forest, One-Class SVM, LOF and  
 326 DBSCAN are employed. Evaluation metrics, including precision, recall, F1 score in Table 2 are  
 327 utilized to assess the performance of the proposed methods. However, clustering methods report  
 328 many anomalies; DBSCAN, If eps is too small, leads to many points treated as noise. LOF also  
 329 depends heavily on n\_neighbors, and Isolation Forest depends on contamination parameter. That is  
 330 why the study further introduces a statistically-based scoring method to identify anomalous nodes.  
 331 The scoring function employs different z-score thresholds of 1.0, 1.5, and 2.0 (95% confidence level).  
 332 Furthermore, we compare the results obtained from our scoring method with the ground truth on  
 333 etherscan.io, providing a case-by-case evaluation of detected time-sensitive anomalies in Table 3.

334 We further compare the TRW-GCN model against the state-of-the-art TGAT method. TGAT is  
 335 specifically designed to incorporate temporal information through time-aware positional encodings  
 336 and attention mechanisms. However, in practice, our experiments revealed significant computational  
 337 and performance challenges when applying TGAT, particularly in complex, high-frequency networks  
 338 such as Ethereum. TGAT’s multi-head attention mechanism introduces substantial overhead due to  
 339 repeated matrix multiplications and attention score computations. Additionally, its dependency on  
 340 fine-grained temporal edge attributes adds complexity to both preprocessing and model execution,  
 341 resulting in long training time and memory inefficiency. In contrast, TRW-GCN’s use of temporal  
 342 random walks allows it to construct meaningful local temporal subgraphs with controlled depth and  
 343 temporal relevance, making it significantly more scalable without sacrificing temporal fidelity. From  
 344 a performance standpoint, TGAT achieved an accuracy of 85.3% detecting 23 anomalies, while our  
 345 TRW-GCN model — coupled with a scoring classifier — has achieved an average accuracy of 94.5%  
 346 detecting 20 anomalies, see Table 4. One likely factor behind this discrepancy is TGAT’s sensitivity  
 347 to the temporal quality and distribution of data. In Ethereum, where transactions are bursty and user  
 348 behavior is non-uniform, TGAT struggles to generalize effectively. Moreover, TGAT’s reliance on  
 349 explicit node identities (e.g., blockchain addresses) complicates indexing and neighborhood retrieval,  
 350 especially in networks with millions of ephemeral or sparsely active nodes. TRW-GCN, in contrast,  
 351 is more robust in such settings due to its walk-based sampling, which implicitly encodes temporal  
 352 structure without depending on densely connected or temporally smooth interactions.

## 353 4 Conclusion

354 The evolution and complexity of the Ethereum network has heightened the urgency for temporal  
 355 anomaly detection methods. Through our research, we’ve demonstrated that the combined TRW-GCN  
 356 method offers a solution to this challenge. This fusion has enabled us to delve deeper into the intricate  
 357 spatial-temporal patterns of Ethereum transactions, offering a refined lens for anomaly detection.  
 358 We have shown the methodology usefulness by expressing and proving three distinct theorems,  
 359 full empirical analysis and evaluation. While this approach is used to obtain the embedding, we  
 360 have compared different clustering and scoring classification methods to obtain highest precision in  
 361 anomaly detection, and verified with the ground truth found on etherscan.io. Furthermore, we have  
 362 demonstrated that the TRW-GCN method improves anomaly detection versus the state-of-the-art  
 363 TGAT method, also proved how probabilistic sampling improves GCN performance in Appendix B.

## References

- [1] B. Yu, M. Li, J. Zhang, and Z. Zhu, “3d graph convolutional networks with temporal graphs: A spatial information free framework for traffic forecasting,” 2019.
- [2] A. Sankar, Y. Wu, L. Gou, W. Zhang, and H. Yang, “Deep neural representation learning on dynamic graphs via self-attention networks,” in *WSDM*, 2020, pp. 519–527.
- [3] A. Pareja, G. Domeniconi, J. Chen, T. Ma, T. Suzumura, H. Kanezashi, T. Kaler, and C. Leiserson, “Evolvegcn: Evolving graph convolutional networks for dynamic graphs,” 2019.
- [4] D. Xu, C. Ruan, E. Korpeoglu, S. Kumar, and K. Achan, “Inductive representation learning on temporal graphs,” 2020, available: <https://openreview.net/forum?id=rJeW1yHYwH>.
- [5] R. Trivedi, M. Farajtabar, P. Biswal, and H. Zha, “Dyrep: Learning representations over dynamic graphs,” 2019, available: <https://openreview.net/forum?id=HyePrhR5KX>.
- [6] Y. Ma, Z. Guo, Z. Ren, E. Zhao, J. Tang, and D. Yin, “Streaming graph neural networks,” 2018.
- [7] G. Nguyen, J. Lee, R. Rossi, N. Ahmed, E. Koh, and S. Kim, “Dynamic network embeddings: From random walks to temporal random walks,” in *2018 IEEE International Conference on Big Data*, 2018, pp. 1085–1092.
- [8] E. Rossi, B. Chamberlain, F. Frasca, D. Eynard, F. Monti, and M. Bronstein, “Temporal graph networks for deep learning on dynamic graphs,” in *International Conference on Learning Representations (ICLR)*, 2020.
- [9] M. Liu, K. Liang, B. Xiao, S. Zhou, W. Tu, Y. Liu, X. Yang, and X. Liu, “Self-supervised temporal graph learning with temporal and structural intensity alignment,” 2023, available: <https://arxiv.org/abs/2302.07491>.
- [10] K. Xu, J. Li, S. Zhang, M. Du, K. Kawarabayashi, and S. Jegelka, “What can neural networks reason about?” 2019.
- [11] S. Li, G. Gou, C. Liu, C. Hou, Z. Li, and G. Xiong, “Ttag: Temporal transaction aggregation graph network for ethereum phishing scams detection,” in *Proceedings of the ACM Web conference 2022 (WWW ’22)*. Association for Computing Machinery, 2022, pp. 661–669.
- [12] J. Wu, Q. Yuan, D. Lin, W. You, W. Chen, C. Chen, and Z. Zheng, “Who are the phishers? phishing scam detection on ethereum via network embedding,” 2020, available: <https://arxiv.org/abs/1911.09259>.
- [13] D. Ofori-Boateng, S. Dominguez, C. Akcora, M. Kantarcioglu, and Y. Gel, “Topological anomaly detection in dynamic multilayer blockchain networks,” in *ECML PKDD 2021: Machine Learning and Knowledge Discovery in Databases*, 2021, pp. 788–804.
- [14] L. Zheng, Z. Li, J. Li, Z. Li, and J. Gao, “Addgraph: Anomaly detection in dynamic graph using attention-based temporal gcn,” in *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence (IJCAI-19)*. International Joint Conferences on Artificial Intelligence Organization, 2019, pp. 1424–1430.
- [15] P. Daian and et al., “Flash boys 2.0: Frontrunning, transaction reordering, and consensus instability in decentralized exchanges,” in *IEEE symposium on security and privacy (SP)*, 2020, pp. 910–927.
- [16] W. Zhang, L. Wei, S. Cheung, Y. Liu, S. Li, L. Liu, and M. Lyu, “Combatting front-running in smart contracts: Attack mining, benchmark construction and vulnerability detector evaluation,” 2023, available: <https://arxiv.org/pdf/2212.12110>.
- [17] Z. Feng, Y. Li, and X. Ma, “Blockchain-oriented approach for detecting cyber-attack transactions,” *Financial Innovation*, vol. 9, no. 81, 2023.
- [18] S. Behfar, R. Mortier, and J. Crowcroft, “Analysis of information propagation in ethereum network using combined graph attention network and reinforcement learning to optimize network efficiency and scalability,” in *Proceedings of the 5th Workshop on Machine Learning and Systems (EuroMLSys ’25)*. ACM, 2025, pp. 238–245. [Online]. Available: <https://doi.org/10.1145/3721146.3721965>

## A Significant weight cancellation

**Theorem 2:** Let  $\mathbb{R}^m$  be a vector space, and let  $h_n \in \mathbb{R}^m$  represent a feature vector. Define a temporal transformation matrix  $T_{n,j} \in \mathbb{R}^{m \times m}$ , where each entry  $t_{ij}$  encodes the temporal weights. Let  $h'_n = T_{n,j}h_n$  be the transformed feature vector.

If the transformation matrix  $T_{n,j}$  exhibits symmetric or complementary weight patterns that cause significant weight cancellation, the difference between the transformed and original vectors,  $\|h'_n - h_n\|_2$ , will be insufficient to surpass a given anomaly detection threshold  $\delta > 0$ . Specifically, weight cancellation occurs if:

$$\sum_{j=1}^m t_{ij}h_{nj} \approx h_{ni}, \quad \forall i \in \{1, 2, \dots, m\}. \quad (10)$$

**Proof:**

419 The transformed feature vector  $h'_n = T_{nj}h_n$  can be expressed component-wise as:

$$h'_{ni} = \sum_{j=1}^m t_{ij}h_{nj}, \quad \forall i \in \{1, 2, \dots, m\}. \quad (11)$$

420 The Euclidean norm of the difference between the transformed and original feature vectors is given  
421 by:

$$\|h'_n - h_n\|_2 = \sqrt{\sum_{i=1}^m \left( \sum_{j=1}^m t_{ij}h_{nj} - h_{ni} \right)^2}. \quad (12)$$

422 For  $\|h'_n - h_n\|_2 > \delta$ , the inequality must hold:

$$\sum_{i=1}^m \left( \sum_{j=1}^m t_{ij}h_{nj} - h_{ni} \right)^2 > \delta^2. \quad (13)$$

423 This implies that, for at least one  $i$ , the inner term  $\sum_{j=1}^m t_{ij}h_{nj} - h_{ni}$  must be at least  $\delta^2$ . Therefore,  
424  $T_{nj}$  must introduce a significant alteration to the distribution of  $h_n$ . Weight cancellation occurs when  
425  $T_{nj}$  has structural properties that lead to minimal change in  $h_n$ . Consider the following cases:

426 - **Symmetry in  $T_{nj}$** : If  $T_{nj}$  is symmetric ( $t_{ij} = t_{ji}$ ) and  $h_n$  has symmetric properties, the transforma-  
427 tion may yield:

$$\sum_{j=1}^m t_{ij}h_{nj} \approx h_{ni}, \quad \forall i. \quad (14)$$

428 In this scenario, the transformed feature vector  $h'_n$  closely resembles original vector  $h_n$ , leading to

$$\|h'_n - h_n\|_2 \approx 0. \quad (15)$$

429 - **Complementary Weights**: If  $T_{nj}$  contains complementary weights, such that certain entries  $t_{ij}$   
430 and  $t_{ik}$  satisfy  $t_{ij} + t_{ik} = 0$ , and if  $h_{nj} \approx h_{nk}$ , then the contributions from  $h_{nj}$  and  $h_{nk}$  cancel each  
431 other out:

$$\sum_{j=1}^m t_{ij}h_{nj} \approx 0, \quad \text{for certain } i. \quad (16)$$

432 - **Spectral Properties of  $T_{nj}$** : If  $T_{nj}$  has eigenvalues close to 1, it behaves similarly to an identity  
433 matrix, resulting in  $h'_n \approx h_n$ . Orthogonality in rows or columns of  $T_{nj}$  may also preserve the  
434 magnitude of  $h_n$ , leading to minimal changes in  $h'_n$ .

435 In scenarios where weight cancellation occurs, the transformation  $T_{nj}$  fails to introduce meaningful  
436 changes to the feature vector  $h_n$ . Consequently, anomalies influenced by temporal factors may not be  
437 detectable, as the difference  $\|h'_n - h_n\|_2$  remains below the threshold  $\delta$ .

## 438 B Improvement of GCN performance with probabilistic sampling

439 **Theorem 3:** Improvement of GCN performance with probabilistic sampling in the context of random  
440 walk sampling.

441 Consider a simplified Ethereum transaction graph with N accounts (nodes), and M transactions (edges)  
442 between them. Prove the performance improvement of a GCN in terms of loss, using probabilistic  
443 sampling for the task of predicting account behaviors, considering the following assumptions:

- 444 1. Nodes (accounts) have features represented by vectors in a feature matrix X.
- 445 2. The adjacency matrix A represents transaction relationships between accounts.
- 446 3. Binary labels Y indicate specific account behaviors.

447 **Proof.**

448 **B.1 Traditional GCN performance**

449 Start with the definition of the normalized graph Laplacian  $L = I - D^{-\frac{1}{2}}AD^{-\frac{1}{2}}$ , where  $D$  is the  
450 diagonal degree matrix and  $A$  is the adjacency matrix.

451 Derive the eigenvalues and eigenvectors of the Laplacian matrix  $L$  and show their significance in  
452 capturing graph structure. Derive the performance of a GCN trained on the full graph using these  
453 eigenvalues and eigenvectors:

454 Step 1: Deriving Eigenvalues and Eigenvectors of the Laplacian matrix  $L$

455 Given the normalized graph Laplacian matrix  $L$ , let  $\lambda$  be an eigenvalue of  $L$  and  $v$  be the corresponding  
456 eigenvector. In the equation  $Lv = \lambda v$ , solving for  $\lambda$  and  $v$ , we get:

$$457 \quad D^{-\frac{1}{2}}AD^{-\frac{1}{2}}v = (1 - \lambda)v \quad (17)$$

$$AD^{-\frac{1}{2}}v = (1 - \lambda)D^{\frac{1}{2}}v \quad (18)$$

458 This equation implies that  $D^{-\frac{1}{2}}AD^{-\frac{1}{2}}$  is a symmetric matrix that is diagonalized by the eigenvectors  
459  $v$  with corresponding eigenvalues  $1 - \lambda$ . The eigenvectors and eigenvalues of  $L$  capture the graph’s  
460 structural information. Larger eigenvalues correspond to well-connected clusters of nodes in the  
461 graph, while smaller eigenvalues correspond to isolated groups or individual nodes.

462 Step 2: Deriving GCN performance using eigenvalues and eigenvectors

463 Now let’s consider a scenario where we’re using a GCN to predict node labels (such as predicting  
464 high-value transactions) on the full graph. The GCN’s propagation rule can be written as:

$$h^{(l+1)} = f(\hat{A}h^{(l)}W^{(l)}) \quad (19)$$

465 where  $h^{(l)}$  is the node embedding matrix at layer  $l$ ,  $f$  is an activation function, and  $\hat{A} = D^{-\frac{1}{2}}AD^{-\frac{1}{2}}$ .  
466 is the symmetrically normalized adjacency matrix, and  $W^{(l)}$  is the weight matrix at layer  $l$ . The key  
467 insight is that if we stack multiple GCN layers, the propagation rule becomes:

$$\begin{aligned} h^{(L)} &= f(\hat{A}h^{(L-1)}W^{(L-1)}) \\ &= f(\hat{A}f(\hat{A}h^{(L-2)}W^{(L-2)})W^{(L-1)}) \dots \end{aligned} \quad (20)$$

468 We can simplify this as:

$$h^{(L)} = f\left(\hat{A}^{(L)}h^{(0)}W^{(0)}\prod_{l=1}^{L-1}W^{(l)}\right) \quad (21)$$

469 Using the spectral graph theory, we know that  $\hat{A}^{(L)}$  captures information about the graph’s structure  
470 up to  $L$ -length paths. The eigenvalues and eigenvectors of  $\hat{A}^{(L)}$  indicate the influence of different  
471 sampled-graphs of length  $L$  on the node embeddings.

472 **B.2 Probabilistic Sampling Approach**

473 In this step, we’ll introduce a probabilistic sampling strategy to select a subset of nodes and their  
474 associated transactions. This strategy aims to prioritize nodes with certain characteristics or properties,  
475 such as high transaction activity or potential involvement in high-value transactions. Assign a  
476 probability  $p_i$  to each node  $i$  based on certain characteristics. For example, nodes with higher  
477 transaction activity, larger balances, or more connections might be assigned higher probabilities.  
478 For each node  $i$ , perform a random sampling with probability  $p_i$  to determine whether the node is  
479 included in the sampled subset. Consider a graph with  $N$  nodes represented as  $N = \{1, 2, \dots, N\}$ .  
480 Each node  $i$  has associated characteristics described by a feature vector  $\mathbf{X}_i = [X_{i,1}, X_{i,2}, \dots, X_{i,k}]$ ,  
481 where  $K$  is the number of characteristics. Define the probability  $p_i$  for node  $i$  as a function of its  
482 feature vector  $\mathbf{X}_i$ :  $p_i = f(\mathbf{X}_i)$ . Here,  $f(\cdot)$  is a function that captures how the characteristics of node  
483  $i$  are transformed into a probability. The specific form of  $f(\cdot)$  depends on the characteristics and the  
484 desired probabilistic behavior. For example,  $f(\mathbf{X}_i)$  could be defined as a linear combination of the  
485 elements in  $\mathbf{X}_i$ :

$$p_i = \sum_{j=1}^K \omega_j X_{i,j} \quad (22)$$

486 Where  $\omega_j$  are weights associated with each characteristic. The weights  $\omega_j$  can be used to emphasize  
 487 or downplay the importance of specific characteristics in determining the probability. After obtaining  
 488  $p_i$  values for all nodes, normalize them to ensure they sum up to 1. Nodes with higher normalized  
 489 probabilities are more likely to be included in the sampled subset.

$$p_{\text{normalized}} = \frac{p_i}{\sum_{j=1}^N p_j} \quad (23)$$

### 490 B.3 Graph Laplacian for Sampled Graph

491 Given the sampled adjacency matrix  $\hat{A}_{\text{sampled}}$ , we want to derive the graph Laplacian  $\hat{L}_{\text{sampled}}$  for the  
 492 sampled graph. The graph Laplacian  $\hat{L}_{\text{sampled}}$  is given by:

$$\hat{L}_{\text{sampled}} = I - \hat{D}_{\text{sampled}}^{-\frac{1}{2}} \hat{A}_{\text{sampled}} \hat{D}_{\text{sampled}}^{-\frac{1}{2}} \quad (24)$$

493 Where  $\hat{D}_{\text{sampled}}$  is the diagonal degree matrix of the sampled graph, where each entry  $d_{ii}$  corresponds  
 494 to the degree of node  $i$  in the sampled graph, and  $\hat{A}_{\text{sampled}}$  is the sampled adjacency matrix.

$$d_{ii} = \sum_{j=1}^{N_{\text{sampled}}} \hat{A}_{\text{sampled},ij} \quad (25)$$

495 The modified Laplacian captures the structural properties of the sampled graph and is essential for  
 496 understanding its graph-based properties. As eigenvalues of the sampled graph, we derive

$$\hat{L}_{\text{sampled}} = I - \hat{D}_{\text{sampled}}^{-\frac{1}{2}} \hat{A}_{\text{sampled}} \hat{D}_{\text{sampled}}^{-\frac{1}{2}} \quad (26)$$

497 as the normalized graph Laplacian for the sampled graph. Let  $\hat{\lambda}_i$  be the  $i$ -th eigenvalue of  $\hat{L}_{\text{sampled}}$   
 498 and  $\hat{v}_i$  be the corresponding eigenvector. We have

$$\hat{L}_{\text{sampled}} \hat{v}_i = \hat{\lambda}_i \hat{v}_i \quad (27)$$

499 The goal is to compare the eigenvalues of  $L$  with the eigenvalues of  $\hat{L}_{\text{sampled}}$  and show convergence  
 500 under certain conditions. As the sample size  $N_{\text{sampled}}$  approaches the total number of nodes  $N$  in the  
 501 original graph,  $\hat{L}_{\text{sampled}}$  converges to  $L$ . Eigenvalues of  $\hat{L}_{\text{sampled}}$  converge to the eigenvalues of  $L$ .

### 502 B.4 Impact on GCN Performance

503 To demonstrate that the performance  $E_{\text{sampled}}$  of a GCN on a sampled graph, is greater than or equal  
 504 to the performance  $E_{\text{full}}$  on the full graph, we use two approaches:

#### 505 1. Reduction of Noise and Retention of Structural Information

506 The total loss  $\mathcal{L}$  of a GCN can be expressed as:

$$\mathcal{L}(h) = \mathcal{L}_{\text{train}}(h) + \mathcal{E}(h) \quad (28)$$

507 where:

- 508 •  $\mathcal{L}_{\text{train}}(h)$ : Loss on the training set.
- 509 •  $\mathcal{E}(h)$ : Generalization error (e.g., noise or overfitting effects).

510 For the sampled graph  $G_{\text{sampled}}$ , the loss becomes:

$$\mathcal{L}(h_{\text{sampled}}) = \mathcal{L}_{\text{train}}(h_{\text{sampled}}) + \mathcal{E}(h_{\text{sampled}}) \quad (29)$$

511 Probabilistic sampling prioritizes nodes with higher relevance (e.g., higher degree or centrality) by  
 512 assigning sampling probabilities  $p_i$ :

$$p_i = f(X_i), \quad p_{\text{normalized}} = \frac{p_i}{\sum_j p_j} \quad (30)$$

513 where  $X_i$  represents node features. By emphasizing relevant nodes, noise is reduced, and:

$$\mathcal{E}(h_{\text{sampled}}) < \mathcal{E}(h) \tag{31}$$

514 Thus, the total loss on the sampled graph satisfies:

$$\mathcal{L}(h_{\text{sampled}}) < \mathcal{L}(h) \tag{32}$$

515 **2. Reduction in Computational Complexity and Faster Convergence**

516 The computational complexity of a GCN is:

$$\mathcal{O}(L \cdot (N + M) \cdot d^2) \tag{33}$$

517 where  $N$  is the number of nodes,  $M$  is the number of edges,  $L$  is the number of layers, and  $d$  is the  
 518 embedding dimension. For the sampled graph  $G_{\text{sampled}}$ , the complexity reduces to:

$$\mathcal{O}(L \cdot (N_{\text{sampled}} + M_{\text{sampled}}) \cdot d^2) \tag{34}$$

519 Since  $N_{\text{sampled}} \ll N$  and  $M_{\text{sampled}} \ll M$ , the sampled graph enables faster convergence. Let the  
 520 convergence rate  $R$  be inversely proportional to the size of the graph:

$$R(G_{\text{sampled}}) > R(G) \tag{35}$$

521 Thus, the sampled graph converges faster and reaches a better minimum of the loss function:

$$\mathcal{L}(h_{\text{sampled}}) \text{ decreases faster compared to } \mathcal{L}(h) \tag{36}$$

522 Given the reduced noise, retention of structural information, and faster convergence, probabilistic  
 523 sampling ensures that:

$$E_{\text{sampled}} > E_{\text{full}} \tag{37}$$

524 **B.5 How TRW impacts on GCN performance as compared to traditional sampling**

525 Let’s delve into empirical justification on why TRW sampling could enhance the performance of  
 526 GCNs, especially in temporal networks like Ethereum. For a detailed mathematical proof on the  
 527 probabilistic sampling in GCN, you are invited to read appendix B1-B4. One issue with traditional  
 528 random walks is the potential for creating "jumps" between temporally distant nodes, breaking the  
 529 temporal consistency. GCNs rely on the local aggregation of information, and since TRW promotes  
 530 smoother temporal signals, GCNs can potentially learn better node representations. Temporal  
 531 consistency ensures that the sequences are logically and temporally ordered. This can be crucial for  
 532 predicting future events or understanding time-evolving patterns, making GCNs more reliable. We  
 533 compare different GCN models (including graphSAGE and graph attention network GAT model) for  
 534 fullgraph, and sampled-graph with traditional and temporal random walk in Figure 4. Although one  
 535 sees little difference between the accuracy of the fullgraph and the sampled-graph in graphSAGE and  
 536 GAT models (18), one can see that traditional random walk and temporal random walk improve GCN  
 537 accuracy, where TRW shows even further improvement than the traditional random walk.

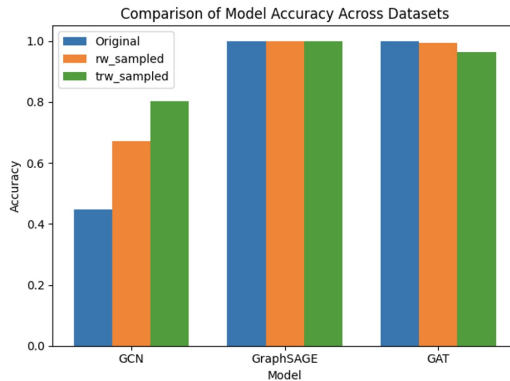


Figure 4: Comparison of fullgraph, traditional RW and TRW-based on sampled graph in 100 blocks.

538 **C Checklist Responses**

- 539 1. **Claims:** *Yes.* The abstract and introduction reflect the contributions of the paper. TRW-GCN  
540 is proposed as a domain-specific temporal GCN variant tailored to Ethereum transaction  
541 networks. The use of probabilistic temporal walks and their effect on anomaly detection are  
542 experimentally demonstrated. The paper acknowledges that the model is not intended as a  
543 general-purpose method, and the scope is clearly limited to complex blockchain structures.
- 544 2. **Limitations:** *Yes.* Limitations are discussed in the text. Notably, the model is tailored to  
545 Ethereum-like graphs and may not generalize to all temporal graph domains. Limitations  
546 in comparison scope (e.g., AddGraph, TGAT) and reliance on temporal features that may  
547 be noisy are acknowledged. We also observed that TGAT results in higher computational  
548 costs, primarily due to its multi-head attention mechanism, which involves multiple passes  
549 of matrix multiplications and attention score computations. Furthermore, TGAT’s reliance  
550 on temporal edge attributes added another layer of complexity, further increasing the  
551 computational burden.
- 552 3. **Theory, Assumptions and Proofs:** *Yes.* We provided all theoretical claims, stated all  
553 assumptions clearly before theorem statements, and provided formal proofs either in the  
554 main paper or appendix.
- 555 4. **Experimental Result Reproducibility:** *Yes.* Both data and code are attached at submission  
556 which also explains how to obtain the paper results.
- 557 5. **Open Access to Data and Code:** *Yes.* Full Ethereum dataset is publicly avail-  
558 able; nevertheless, we provide our created dataset and the code in the github link  
559 <https://github.com/stefankam/temporal-spacial-anomaly-detection>, which is anonymized.
- 560 6. **Experimental Setting/Details:** *Yes.* Full training and testing splits, model hyperparameters,  
561 walk lengths, and walk counts are provided in the text. Comparisons with TGAT and other  
562 unsupervised methods (e.g., SVM, ISOForest) are described.
- 563 7. **Experiment Statistical Significance:** *Yes.* The scoring method is based on z-score thresh-  
564 olds (1.0, 1.5, 2.0), corresponding to standard confidence levels (e.g., 95%). The reported  
565 precision/recall/f1 are averaged over multiple thresholds and visualized. Confidence intervals  
566 are also included in Figure 1.
- 567 8. **Experiments Compute Resource:** *Yes.* Experiments were run on our department server  
568 running Linux equipped with a single GPU (NVIDIA A100 80GB PCIe), and 251Gi RAM..
- 569 9. **Code of Ethics:** *Yes.* The research conforms to NeurIPS Code of Ethics. No human or  
570 sensitive data was used. All datasets are public and open.
- 571 10. **Broader Impacts:** *Yes.* The paper discusses anomaly detection and classification systems.  
572 Limitations of false positives are acknowledged specially in the clustering methods like  
573 dbscan which demonstrate high number of anomalies’ detection. Future work could help  
574 mitigate misclassification risks, and further automation.
- 575 11. **Safeguards:** *N/A.* No pretrained models with dual-use risks are released. The framework is  
576 domain-specific and does not apply to general-purpose generative tasks.
- 577 12. **Licenses:** *Yes.* Ethereum transaction data is public and under open access. All reused  
578 datasets (e.g., etherscan.io) are cited appropriately. Libraries used include PyTorch Geomet-  
579 ric (MIT License).
- 580 13. **Assets:** *No.* While no new datasets are introduced, the model artifacts and scripts will be  
581 documented and released.
- 582 14. **Crowdsourcing and Research with Human Subjects:** *N/A.* No human data or crowdsourc-  
583 ing was involved.
- 584 15. **IRB Approvals:** *N/A.* Not applicable as no human or user-generated content was analyzed.
- 585 16. **Declaration of LLM Usage:** *Yes.* LLMs (e.g., ChatGPT) were used only for editing  
586 and understanding of some technical concepts. They did not influence model design or  
587 methodology.