# A Method for Identifying Anomalous Bitcoin Transactions Based on Image-based Representations of Transaction Features

**Zilin Xu**
*School of Computer Science, Wuhan University, Wuhan, Hubei, China*

**Abstract: Cryptocurrencies such as Bitcoin are widely used in financial transactions and cross-border payments due to their decentralized and anonymous nature; however, they also provide a covert channel for anomalous transactions such as money laundering and fraud. Addressing the limitations of traditional abnormal transaction detection methods in utilizing transaction network structures and node attribute information, this paper proposes an abnormal transaction detection method based on Node attribute heatmaps.The method first constructs a directed graph using the Elliptic Bitcoin dataset and extracts local k-hop subgraphs centered on labeled transaction nodes. Subsequently, high-dimensional node attribute features are mapped to node attribute heatmaps, transforming the graph-based anomalous transaction detection problem into an image classification task. Finally, a convolutional neural network is employed to perform binary classification on the image samples. To validate the method's effectiveness, a comparative experiment with a Graph Convolutional Network (GCN) was designed.Experimental results show that the Node attribute heatmap method outperforms both GCN and traditional adjacency matrix image methods in metrics such as Accuracy, Recall, and F1-score, indicating that this method can more fully capture the behavioral characteristics of transaction nodes, thereby enhancing the performance of anomalous transaction detection.**

**Keywords: Anomalous Transaction Detection; Bitcoin; Graph Representation Learning; Node Attribute Heatmap; Convolutional Neural Network (CNN)**

## 1. Introduction

With the rapid development of blockchain technology and digital currencies, cryptocurrencies such as Bitcoin have been widely adopted in the fields of financial payments and asset trading due to their decentralization, anonymity, and convenience for cross-border circulation. While these technical characteristics provide freedom and convenience in transactions, they also result in a lack of centralized regulation over transaction activities found in traditional financial systems. Coupled with the anonymity and concealment of cryptocurrency transactions, this has gradually made them an important tool for money laundering, fraud, illegal financing, and other financial crimes.The Bitcoin transaction network inherently exhibits graph-structured characteristics, with complex network topologies formed by fund flow relationships between transaction nodes. In this context, nodes correspond to transaction entities, and edges represent the paths of fund transfers. Within such networks, anomalous transactions often concentrate in localized structures or behavioral patterns of specific nodes. Identifying potential anomalous nodes within vast and dynamic transaction graphs has become a core challenge in Anti-Money Laundering (AML) research.

Current research on anomaly detection for graph-structured data primarily focuses on Graph Neural Networks (GNNs) and Graph-to-Image methods. GNN methods achieve node representation learning by aggregating features from a node's neighbors and perform well in node classification and graph classification tasks. However, they suffer from issues such as complex model structures, high training and inference costs, insufficient interpretability, and limited ability to visualize node attributes.Graph-to-Image methods achieve graph classification by mapping graph structures to images and using convolutional neural networks to extract features. However, most existing methods rely on adjacency matrices, retaining only topological connection information, which results in sparse generated images that struggle to fully capture the high-dimensional attributes of nodes. Yet, transaction amounts, behavioral statistics, and time-related features of nodes play a critical role in detecting anomalous transactions.

To address these issues, this paper proposes an anomalous transaction detection method based on Node attribute heatmaps.The study first constructs a directed graph network of Bitcoin transactions using the Elliptic dataset and extracts local subgraphs within the k-hop neighborhood of labeled transaction nodes to model the local structure of the transaction network. Based on this, the high-dimensional attributes of nodes are organized into a two-dimensional feature matrix, and a two-dimensional image is generated via heatmap mapping, allowing the numerical variations of node attributes to be intuitively represented in visual space through color intensity and texture.This representation method simultaneously preserves graph structural information and node behavioral features, providing rich input for Graph Neural Networks (GNN). Experimental results demonstrate that node attribute heatmaps offer significant advantages in capturing transaction behavior patterns and enhancing the ability to identify anomalous transactions. Overall, by unifying the representation of graph structure and node attribute features, this method not only improves the accuracy of anomalous transaction detection but also provides a viable technical approach for the image-based representation of financial transaction graphs.

## 2. Literature Review

In recent years, research on financial anomaly detection has primarily focused on two categories of methods: traditional machine learning methods and graph-based learning methods. Traditional methods typically rely on manually designed statistical features and combine them with logistic regression, random forests, or support vector machines to perform transaction classification. However, on large-scale transaction networks, such methods struggle to fully capture the complex structural relationships between nodes, thereby limiting their ability to identify anomalous transaction patterns. To address this issue, researchers have gradually introduced graph-based learning methods into the anomaly detection task.The Elliptic Bitcoin dataset proposed by Weber et al. provides 203,769 transaction nodes, 234,355 transaction edges, and 166-dimensional node attribute features containing both local features and global aggregate features, offering a standardized experimental platform for financial anomalous transaction detection research [1].Building on this foundation, Graph Neural Network (GNN) methods such as Graph Convolutional Networks (GCN) [2] and Graph Attention Networks (GAT) [3] have been widely adopted. These methods learn node representations from graph structures by aggregating neighborhood information and have achieved good performance in detecting anomalous transactions.The Graph Isomorphism Network (GIN) proposed by Xu et al. [4] further enhances the model's ability to distinguish isomorphism among different graph structures by adjusting weights, thereby improving structural representation capabilities for graph classification tasks. However, these graph learning methods typically suffer from issues such as complex model parameters, high training costs on large-scale dynamic graphs, and insufficient interpretability, and they remain inadequate in visualizing high-dimensional node attributes [5].

Although Graph Neural Network (GNN) methods can capture local features, they often face performance bottlenecks caused by over-smoothing when processing large-scale networks containing complex financial features, and they lack intuitive visualization methods to display node behavioral characteristics [5].To address this issue, the Graph-to-Image method was proposed. Its core idea is to map non-Euclidean graph structures to regular image matrices through specific rules, thereby leveraging the powerful Local features capabilities of Convolutional Neural Networks (CNNs) to perform classification tasks [6]. This method enables end-to-end learning using mature computer vision models and provides analytical tools such as feature map visualization.In prior research, Niepert et al. linearized local graph structures into image-like input sequences by selecting central nodes and normalizing their neighborhoods [6]; meanwhile, Hegde et al. explored subgraph classification methods based on Adjacency matrix image representations, utilizing matrix texture information to identify network features [7]. Drawing on these approaches, this paper proposes a node attribute heatmap representation method.

Although Graph-to-Image methods perform well in graph classification tasks, most existing methods retain only the Topological structure of the graph, and their representation of node attribute features remains limited.In financial anomaly detection tasks, node-level features such as transaction amounts, behavioral statistics, and temporal characteristics play a crucial role in identifying anomalous transactions. Therefore, how to integrate graph structural information with node attribute features into image-based representations to construct efficient and intuitive image representations remains a key issue in current research.

Compared to existing adjacency matrix-based Graph-to-Image methods, the node attribute heatmap representation method proposed in this paper maps high-dimensional node attributes into two-dimensional visual images with rich textural structures, thereby enhancing the learning ability of convolutional neural networks to recognize anomalous trading patterns. Furthermore, by incorporating the construction of local transaction subgraphs, this method not only preserves graph structural information but also fully integrates node attribute features into the image representation, providing a more intuitive and effective means for detecting anomalous transactions.

## 3. Method

### 3.1 Overall Framework

This paper proposes a financial anomalous transaction detection framework based on Node attribute heatmap representation, with its overall workflow illustrated in Figure 1.The study first extracts transaction nodes, transaction edges, and node attribute features from the Elliptic Bitcoin dataset to construct a directed graph network of Bitcoin transactions. After constructing the network, k-hop subgraphs within the neighborhood of a labeled transaction node are extracted, thereby modeling the local structure of the transaction network. For each k-hop subgraph, this paper generates two types of image representations: traditional adjacency matrix images and node attribute heatmaps.The node attribute heatmap maps the high-dimensional attribute matrix of a node into a two-dimensional visual image, enabling the intuitive visualization of transaction node behavior patterns and providing rich input for convolutional neural networks. Finally, these images are fed into a convolutional neural network (CNN) for training to perform the binary classification task of distinguishing between illegal and legal transaction nodes, thereby achieving the identification of anomalous transactions.

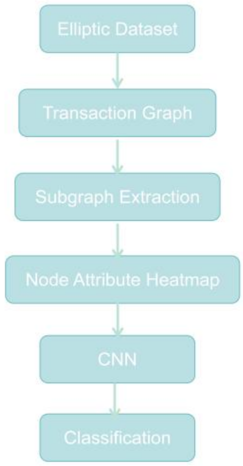


**Figure 1. Overall Framework of This Experiment**

### 3.2 Construction of the Transaction Graph

The transaction relationships in the Elliptic Bitcoin dataset inherently possess graph-structural characteristics. This paper models the entire transaction network as a directed graph $G = (V, E)$, where: $V$ represents the set of transaction nodes, with each node corresponding to a Bitcoin transaction; $E$ represents the set of transaction edges, and a directed edge $e = (u, v)$ indicates a fund transfer relationship between transaction nodes $u$ and $v$.

This paper first reads the transaction edge file elliptic_txs_edgelist.csv from the dataset and constructs the directed transaction graph using NetworkX. This network comprises over 200,000 nodes and 230,000 edges, forming a complex and large-scale transaction structure.Additionally, each transaction node in the dataset contains 166-dimensional attribute features, including local features and aggregate features. These node attribute details can, to a certain extent, reflect transaction behavior patterns, providing a crucial basis for subsequent anomaly detection.

### 3.3 Local Subgraph Extraction

Due to the massive scale of the Bitcoin transaction network, performing learning directly on the complete graph incurs high computational costs. At the same time, anomalous transaction behavior is typically concentrated within local transaction relationships; therefore, this paper adopts a local subgraph modeling strategy. Specifically, using labeled transaction nodes as central nodes, local transaction subgraphs are extracted via the k-hop subgraph method.For a central node v, its local subgraph is defined as $G_v=(V_v,E_v)$, where $V_v$ denotes the set of nodes no more than k hops away from the central node, and $E_v$ denotes the corresponding set of edges.

In implementation, this paper utilizes the ego_graph method provided by NetworkX to extract local subgraphs, and sets k=3 to construct a three-hop neighborhood subgraph.To avoid generating subgraphs that are too small and thus lack sufficient feature information, while ensuring the validity of the data samples, a minimum node count threshold is further established to filter out subgraphs with too few nodes. Through this local subgraph extraction strategy, structural information surrounding the central trading nodes can be effectively preserved, providing a rich feature foundation for subsequent image representation and the learning of anomalous transaction patterns.

### 3.4 Adjacency Matrix Image Representation

To investigate the performance of different image representation methods in the task of detecting anomalous transactions, this paper first constructs image representations using the traditional Adjacency matrix approach.For a local subgraph Gv, its adjacency matrix is defined as $A \in \mathbb{R}^{N \times N}$, where N denotes the number of nodes in the subgraph, and the matrix element Aijindicates whether a connection exists between nodes i and j. In practice, this paper uses NetworkX to convert the local subgraph into an adjacency matrix and maps the matrix to a grayscale image as shown in Figure 2 using the imshow function of matplotlib, thereby constructing a Graph-to-Image representation.

This adjacency matrix image preserves the topological structure of the local transaction network; however, because the matrix itself is typically sparse, the generated image contains limited meaningful texture information, making it difficult to fully express the behavioral attributes of transaction nodes. Therefore, although adjacency matrix images can be used to capture network structural patterns, they have significant limitations in expressing node attribute information, providing a direction for improvement in subsequent node attribute heatmaps.

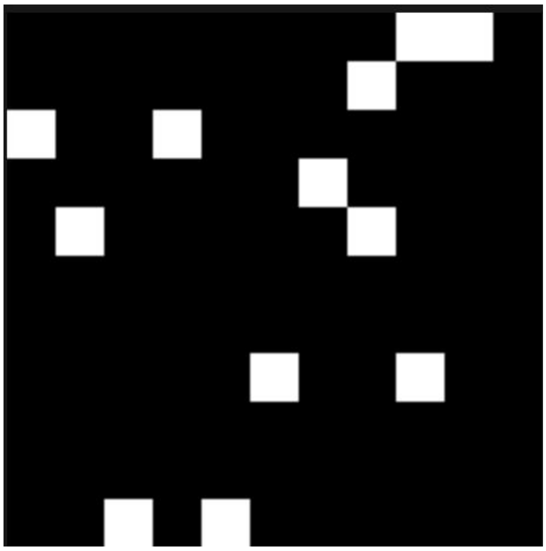

**Figure 2. Example of an Adjacency Matrix Image Representation**

### 3.5 Node Attribute Heatmap Representation

Traditional adjacency matrix image representations retain only the structural connectivity information of the graph, making it difficult to fully reflect the high-dimensional attribute features of trading nodes. However, in financial anomalous transaction detection tasks, node attributes often contain key trading behavior patterns, such as trading statistical features, local aggregate features, and time-related features. To address this, this paper proposes a node attribute heatmap representation method to enhance the visual expression of trading behavior patterns.For a local subgraph $G_v$, assuming it contains N nodes, each with D-dimensional attribute features, a node feature matrix$X \in \mathbb{R}^{N \times D}$ can be constructed, where each row represents a node's attribute vector and each column corresponds to a feature dimension.

After constructing the node feature matrix, this paper normalizes it and uses a heatmap to convert it into a two-dimensional image, as shown in Figure 3. In the heatmap, different numerical values are represented by color intensity, allowing the distribution patterns of node attributes to be intuitively presented as visual textures.Compared to traditional adjacency matrix images, the node attribute heatmap not only preserves node attribute information but also generates richer textural structures, thereby providing ample input features for convolutional neural networks and helping to enhance the ability to identify anomalous transaction patterns.

 

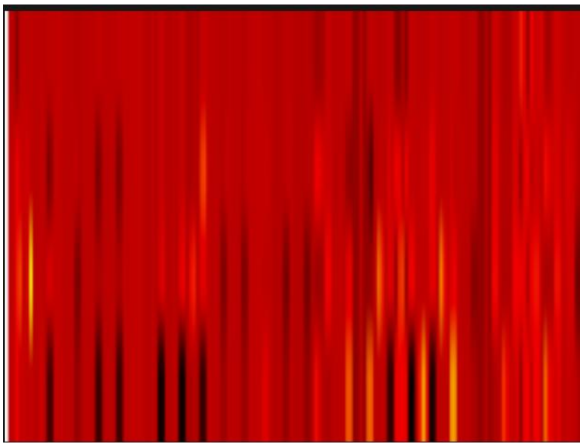

**Figure 3. Example of a Node Attribute Heatmap**

As shown in Figure 3, compared to adjacency matrix images, node attribute heatmaps possess richer textural structures and numerical distribution patterns, enabling a more intuitive representation of high-dimensional attribute information within transaction nodes.

### 3.6 CNN Classification Model

After generating the images, this paper employs a Convolutional Neural Network (CNN) to classify trading nodes for anomalies. Convolutional neural networks possess strong feature extraction capabilities in image classification tasks; they can automatically learn local texture patterns in images through convolution operations, thereby capturing the behavioral characteristics of anomalous transaction nodes. To this end, this paper constructs a simple CNN classification model, whose structure includes convolutional layers, pooling layers, and fully connected layers.The convolutional layer is responsible for extracting local texture features from the images; the pooling layer reduces the feature dimension and computational complexity through dimension reduction operations; and the fully connected layer performs the final binary classification task.

During training, the generated image samples are divided into training and testing sets, and the model is optimized using a cross-entropy loss function. Through model training, the CNN can effectively distinguish between illegal and legitimate transaction nodes, enabling the automatic identification of anomalous transactions.

## 4. Experiments and Analysis

### 4.1 Dataset and Experimental Environment

This study utilizes the Elliptic Bitcoin dataset. This dataset is constructed from the real Bitcoin transaction network and has been widely used in research on financial anomalous transactions. The dataset primarily consists of three components:

(1) Transaction edge file (elliptic_txs_edgelist.csv), which describes the fund flow relationships between transaction nodes;

(2) Node attribute file (elliptic_txs_features.csv), containing 166-dimensional attribute features for each transaction node;

(3) Node label file (elliptic_txs_classes.csv), where label 1 denotes illegal transactions, label 2 denotes legal transactions, and "unknown" denotes an unknown category.

In the experiment, this paper retains only transaction nodes with explicit labels and constructs local transaction subgraphs centered on them. The entire transaction network comprises approximately 200,000 nodes and 230,000 edges, representing a typical large-scale sparse graph structure dataset.

The experimental environment is as follows:

Operating System: Windows 11

Development environment: Python 3.10

Deep learning framework: PyTorch

Graph processing tool: NetworkX

Experimental platform: Anaconda virtual environment

### 4.2 Data Preprocessing and Image Generation

Due to the massive scale of the original transaction network, training directly on the complete graph would incur significant computational overhead. Therefore, this paper employs a local subgraph modeling approach to construct training samples. Using labeled transaction nodes as centers, local transaction subgraphs are extracted via a three-hop neighborhood (3-hop).During the actual construction process, to avoid insufficient information due to excessively small subgraphs, a minimum node count threshold was set to filter out subgraph samples with too few nodes. This strategy effectively preserves the local structural information around the central node, providing a reliable feature foundation for subsequent image representation.

During the image generation phase, this paper develops two types of image representation methods for each local subgraph: adjacency matrix image representation and node attribute heatmap representation. The adjacency matrix image preserves the topological structure information of the local transaction network by converting the local subgraph into an adjacency matrix and generating a two-dimensional image using grayscale mapping.The Node attribute heatmap representation first extracts 166-dimensional attribute features from all nodes in the subgraph to construct a node feature matrix, which is then normalized. Subsequently, a two-dimensional image is generated using a heatmap approach. Upon completion of image generation, each image sample includes a corresponding class label and can be directly used for training convolutional neural networks to identify anomalous transaction nodes.

### 4.3 Model Training Setup

This paper employs a Convolutional Neural Network (CNN) to perform the task of classifying anomalous transaction nodes. The constructed model primarily consists of two convolutional layers, a pooling layer, and a fully connected layer. The convolutional layers are used to extract local texture features from the images, the pooling layer reduces the feature dimension and computational complexity through dimensionality reduction, and the fully connected layer performs the final binary classification task.In the experiments, the generated image samples were divided into training and testing sets in an 8:2 ratio to ensure the adequacy of model training and the reliability of testing.

During training, the model employed the Adam optimizer for parameter updates, with a learning rate set to 0.001 and the cross-entropy loss function selected as the loss function. In the method validation phase, both adjacency matrix images and Node attribute heatmaps were used as inputs for training and testing, respectively. The classification performance under different image representation methods was compared to evaluate the contribution and effectiveness of node attribute information in anomalous transaction detection.

### 4.4 Comparative Experimental Results and Analysis

To comprehensively validate the effectiveness of the method proposed in this paper, the experiment selected an image representation method based solely on the Adjacency matrix and the classic Graph Neural Network (GNN) model GCN (Graph Convolutional Network) as comparison benchmarks. In addition to classification accuracy (Acc), the evaluation metrics specifically included precision (Precision), recall (Recall), and F1-score to comprehensively measure the model's performance in identifying illegal transactions.The comparison data are shown in Table 1.

**Table 1. Experimental Results of Different Image Representation Methods**

| Model Method | Accuracy (Acc) | Precision | Recall | F1-score |
|---|---|---|---|---|
| Adjacency matrix + CNN | 63.44% | 0.5521 | 0.4812 | 0.5142 |
| GCN (Baseline) | 73.08% | 0.9744 | 0.7206 | 0.8285 |
| Proposed Method (Heatmap + CNN) | 83.66% | 0.867 | 0.8743 | 0.8707 |

Based on the experimental data, the GCN model demonstrates an extremely high precision (0.9744), but its recall (0.7206) is significantly low. This performance discrepancy stems from the logic of the GCN's graph convolution operator.Graph convolution aggregates features from neighboring nodes through a Laplacian smoothing mechanism, which tends to produce excessive smoothing effects when processing highly discrete financial attributes, such as those found in the Elliptic dataset. This mechanism makes it difficult for the model to capture deep-seated anomalous behavior patterns, resulting in a large number of false negatives when identifying illegal transactions.

In contrast, our method achieves an F1-score of 0.8707, representing an approximately 5% improvement over GCN. The core improvement lies in the recall rate, where our method achieves 0.8743, which is 15.37% higher than GCN.

The heatmap representation transforms 166-dimensional high-dimensional attributes into two-dimensional spatial textures, enabling the CNN to extract more subtle and complex behavioral pattern features through convolutional kernels.Even under constraints on subgraph size (k=3), the method still demonstrates exceptional feature capture capabilities. This high recall effectively reduces the concealment of money laundering activities, enabling more precise identification of illicit fund flows and aligning with the practical regulatory requirement for "comprehensive coverage" in financial risk control scenarios.

 

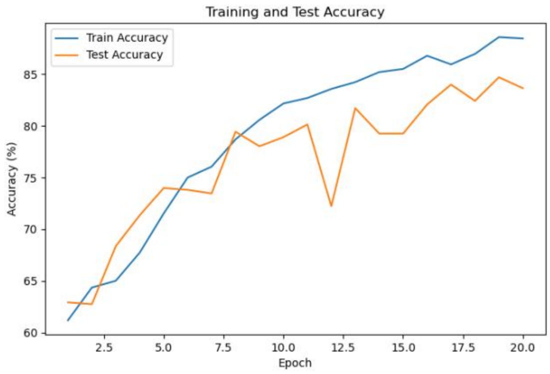


**Figure 4. Model Training Accuracy Curve**

As shown in Figure 4, the model's accuracy on both the training and test sets generally increases with the number of training iterations. During the early stages of training, the model's accuracy improved rapidly, indicating that the convolutional neural network can effectively learn transaction behavior features from heatmaps. As training progressed, the test set accuracy gradually stabilized, ultimately achieving a classification accuracy of 83.66% on the test set.Furthermore, while the accuracy on the training set continued to improve in the later stages of training, the increase in accuracy on the test set slowed, indicating that the model exhibits a certain degree of overfitting.

### 4.5 Analysis of Experimental Results

The experimental results show that, compared to traditional adjacency matrix representations, node attribute heatmap representations demonstrate superior classification performance in the anomalous transaction detection task. This advantage can be understood from several perspectives. On the one hand, adjacency matrix images contain only the connectivity relationships of the graph structure; the generated images are typically sparse and contain limited effective texture information, making it difficult to fully reflect anomalous transaction behavior patterns.On the other hand, node attribute heatmaps can simultaneously preserve high-dimensional attribute information of transaction nodes, including local statistical features and aggregate features, thereby enhancing the representational capacity of anomalous transaction patterns within the image. Furthermore, images generated from heatmaps possess richer textural structures, enabling convolutional neural networks to more easily learn distinguishing features between different categories.

During training, it was observed that while the model's accuracy on the training set continued to improve with increasing training iterations, the accuracy on the test set declined during certain phases, indicating a degree of overfitting. The primary causes of overfitting include the relatively limited number of illegal transaction samples, the still-restricted dataset size, and the large parameter size of the convolutional neural network.

Overall, the node attribute heatmap representation method effectively enhances the expressiveness of transaction node attribute information, providing rich input to the CNN, thereby achieving significant experimental results in the anomalous transaction detection task.

## 5. Conclusions and Outlook

### 5.1 Conclusions

To address the challenge that traditional financial anomaly detection methods struggle to fully leverage node attribute information, this paper proposes a Bitcoin anomalous transaction detection method based on node attribute heatmap representations. In terms of method design, the study first constructs a directed graph using the Elliptic Bitcoin dataset and extracts local k-hop subgraphs centered on labeled transaction nodes, thereby enabling fine-grained modeling of local transaction structures.Based on this, both adjacency matrix image representations and Node attribute heatmaps are generated, transforming the graph structure learning problem into an image classification task. Subsequently, these images are input into a convolutional neural network for training to achieve binary classification between illegal and legitimate transactions.

Experimental results show that the traditional adjacency matrix image representation achieves a classification accuracy of approximately 63.44% on the test set, whereas the Node attribute heatmap representation method achieves 83.66% on the same test set, significantly outperforming the traditional method.These results indicate that, compared to relying solely on topological structure, node attribute information plays a crucial role in financial anomaly detection. By mapping high-dimensional node attribute features into node attribute heatmaps, the expressiveness of transaction behavior patterns is effectively enhanced, thereby improving the convolutional neural network's ability to identify anomalous transactions.

Overall, this paper validates the feasibility of the Graph-to-Image method for detecting anomalous transactions in the financial sector. It also demonstrates that node-attribute-based image representations hold significant research value in this field, providing a reference for future method improvements and applications.

### 5.2 Outlook

Although the Node attribute heatmap representation method has achieved satisfactory results in the anomalous transaction detection task, there is still room for further optimization. Future research can be expanded in multiple directions. First, attempts can be made to integrate various image representation methods.This paper primarily focuses on Node attribute heatmaps, whereas future research could further integrate information such as Adjacency matrices and spectral embeddings to construct multi-channel image representations, thereby enhancing the ability to capture graph structure. Second, Graph Neural Networks (GNN) could be introduced for comparative studies.Current experiments primarily rely on convolutional neural networks (CNNs) for classification tasks; future work could systematically compare this method with Graph Neural Network (GNN) approaches such as GCNs and GATs to evaluate performance differences among various graph representation learning methods in anomaly detection.

Furthermore, enhancing model interpretability is a key direction for future research. Since the method in this paper essentially falls under an image classification framework, it can be combined with interpretability methods such as Grad-CAM to analyze the key regions of interest in the convolutional neural network, thereby improving the interpretability of the results of anomalous transaction detection. Finally, the generalization capability of the method remains to be verified.The experiments in this paper are primarily based on the Elliptic Bitcoin dataset. In the future, the method could be further extended to other financial transaction networks, anti-fraud, and risk propagation analysis tasks to evaluate its applicability and robustness across different financial scenarios.

## References

[1] Weber, M., Domeniconi, G., Chen, J., Weidele, D. K. I., Bellei, C., Robinson, T., & Leiserson, C. E. (2019). Anti-money laundering in bitcoin: Experimenting with graph convolutional networks for financial forensics. arXiv preprint arXiv:1908.02591.

[2] Kipf, T. N., & Welling, M. (2016). Semi-supervised classification with graph convolutional networks. arXiv preprint arXiv:1609.02907.

[3] Veličković, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., & Bengio, Y. (2017). Graph attention networks. arXiv preprint arXiv:1710.10903.

[4] Xu, K., Hu, W., Leskovec, J., & Jegelka, S. (2018). How powerful are graph neural networks?. arXiv preprint arXiv:1810.00826.

[5] Zhou, J., Cui, G., Hu, S., Zhang, Z., Yang, C., Liu, Z., Wang, L., Li, C., & Sun, M. (2020). Graph neural networks: A review of methods and applications. AI Open, 1, 57–81. https://doi.org/10.1016/j.aiopen.2021.01.001

[6] Niepert, M., Ahmed, M., & Kutzkov, K. (2016, June). Learning convolutional neural networks for graphs. In *International conference on machine learning* (pp. 2014–2023). PMLR.

[7] Hegde, K., Magdon-Ismail, M., Ramanathan, R., & Thapa, B. (2018). Network signatures from image representation of adjacency matrices: Deep/transfer learning for subgraph classification. arXiv preprint arXiv:1804.06275.

 