

# Revisiting Graph-Based Fraud Detection in Sight of Heterophily and Spectrum

Fan Xu<sup>3</sup>, Nan Wang<sup>1\*</sup>, Hao Wu<sup>3</sup>, Xuezhi Wen<sup>1</sup>, Xibin Zhao<sup>2\*</sup>, Hai Wan<sup>2</sup>

<sup>1</sup> School of Software, Beijing Jiaotong University, Beijing, China

<sup>2</sup> BNRist, KLISS, School of Software, Tsinghua University, Beijing, China

<sup>3</sup> IAT, University of Science and Technology of China, Hefei, China

{markxu, wuhao2022}@mail.ustc.edu.cn;

{wangnanbjtu, 22126399}@bjtu.edu.cn; {zxb, wanhai}@tsinghua.edu.cn

## Abstract

Graph-based fraud detection (GFD) can be regarded as a challenging semi-supervised node binary classification task. In recent years, Graph Neural Networks (GNN) have been widely applied to GFD, characterizing the anomalous possibility of a node by aggregating neighbor information. However, fraud graphs are inherently heterophilic, thus most of GNNs perform poorly due to their assumption of homophily. In addition, due to the existence of heterophily and class imbalance problem, the existing models do not fully utilize the precious node label information. To address the above issues, this paper proposes a semi-supervised GNN-based fraud detector SEC-GFD. This detector includes a hybrid filtering module and a local environmental constraint module, the two modules are utilized to solve heterophily and label utilization problem respectively. The first module starts from the perspective of the spectral domain, and solves the heterophily problem to a certain extent. Specifically, it divides the spectrum into various mixed-frequency bands based on the correlation between spectrum energy distribution and heterophily. Then in order to make full use of the node label information, a local environmental constraint module is adaptively designed. The comprehensive experimental results on four real-world fraud detection datasets denote that SEC-GFD outperforms other competitive graph-based fraud detectors. We release our code at <https://github.com/Sunxkissed/SEC-GFD>.

## Introduction

With the rapid development of the Internet, various fraudulent entities emerge in fraud activities, and the huge losses caused by this have attracted continuous attention from industry, academia and state agencies (Wang et al. 2018; Pourhabibi et al. 2020; Liu et al. 2021a; Zhang et al. 2021). For instance, fraudulent behaviors aimed at online payment have caused huge property losses to clients; and fraudulent or malicious comments on websites may bring negative effects on merchants and consumers.

In recent years, graph-based fraud detection methods (Hooi et al. 2017; Qin et al. 2022) have been widely used in practical applications. Using entities in fraud scenarios as nodes and interactions between entities as edges, graph-level fraud detection can be achieved. Furthermore, based on the

constructed fraud graph, as a deep neural network for graph representation learning, Graph Neural Network (GNN) (Jin et al. 2021) has been demonstrated to be effective in this task. Specifically, GNN employs neural modules to aggregate neighbor information and update node representations, enabling the discrimination of fraudulent entities.

However, fraud detection scenarios are inherently heterophily (Zhu et al. 2021), *i.e.*, target nodes (like node 5 in Figure 1) and their direct neighbors are prone to different class labels or features. In practical applications, fraudulent entities may use benign entities as springboards to perform high-order interactions with multi-hop neighbors to accomplish fraudulent actions or transmit fraudulent information, *e.g.*, fraud node 3 and 6 are two-hop neighbors in Figure 1. For instance, in financial transaction networks (Akoglu, Tong, and Koutra 2015), fraudsters often use non-fraudulent users to conduct transactions.

Due to the existence of heterophily, the performance of conventional graph neural network is unsatisfactory. This is because vanilla GNN is naturally a low-pass filter, which forces the representation of adjacent nodes to be similar during the information aggregation process. Then it may destroy the discriminative information of anomaly nodes. Furthermore, in the domain of fraud detection, heterophily graph datasets exhibit highly imbalanced class distribution properties. This phenomenon results in anomaly nodes being submerged in normal nodes, which manifests as high homophily of normal nodes and high heterophily of anomaly nodes. This is also called the common camouflage problem (Zhang et al. 2023) in fraud detection scenarios. Moreover, current semi-supervised fraud detection methods based on Graph Neural Networks (GNNs) suffer from limitations in terms of label utilization. These approaches typically learn node representations through a message passing mechanism, extract node embeddings from the final layer of the network, and incorporate label information during the training process. Evidently, these approaches overlook the intricate high-order associations between node label information, node characteristics, and their contextual surroundings.

The current paramount concern entails addressing the heterophily issue within the context of extreme category imbalance. The advanced methods can be divided into two types. The first type of the methods is to design a new message aggregation mechanism (Xu et al. 2023) to simultane-

\*Corresponding author.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

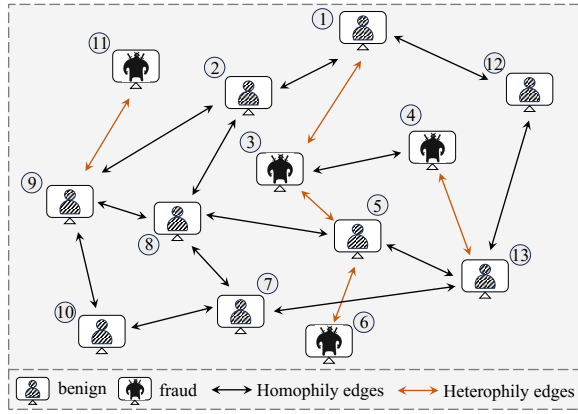


Figure 1: Schematic diagram of fraud relationship in real scenarios.

ously or separately consider the characteristics of homophily and heterophily edges. The second type of the methods is to identify the heterophily edges on fraud graphs and then perform pruning. According to recent researches, scholars found the spectral ‘right-shift’ phenomenon (Tang et al. 2022) on both synthetic and real datasets, proving that the higher the anomaly ratio, the more high-frequency components on the graph spectrum. In addition, researchers also demonstrate that, when using the ground truth labels as signals, the proportion of heterophily edges is positively correlated with the spectral energy distribution (Gao et al. 2023a). Based on these, we conclude that heterophily is not universally harmful, instead it ought to respond distinctively when meeting different filters. In order to verify our conjecture, we conduct sufficient experiments, and the specific related discussions are placed in experiments.

In order to design a GNN-based fraud detector that adapts to high heterophily and has remarkable label utilization, we first propose a conjecture that heterophily edge clipping responds differently to different filters, and then conduct comprehensive experiments to verify the conjecture. Based on the results, we design a **Spectrum-Enhanced and Environment-Constrained Graph Fraud Detector (SEC-GFD)** to comprehensively integrate rich spectrum and label information into a fraud detector. The framework of the proposed method is shown in Figure 2, consisting of two different modules that address the aforementioned two challenges respectively. The hybrid-pass filter module employs a sophisticated message passing technique utilizing the decomposed spectrum to acquire comprehensive high-frequency representations. The local environmental constraint module enhances the higher-order connections between target nodes and their surrounding environments, thereby mitigating the issue of limited label utilization. The main contributions of this method are summarized as follows:

- To address the challenge of heterophily, this paper employs spectrum analysis to partition the graph spectrum into hybrid frequency bands. Then a novel hybrid spectral filter is designed to conduct message passing on these frequency bands separately and finally aggregate.

- With the objective of enhancing label utilization, we introduce a local environmental constraint module to assist the training procedure.
- We conduct extensive experiments on four real-world benchmark datasets, and the results demonstrate the superior effectiveness and robustness of our proposed SEC-GFD over state-of-the-art methods.

## Related Work

**Semi-supervised Node Classification.** Semi-supervised node classification is one of the most important tasks in graph learning. In recent years, GNNs have achieved excellent performance in semi-supervised node classification, like GCN (Kipf and Welling 2016), GraphSAGE (Hamilton, Ying, and Leskovec 2017), GAT (Velivcković et al. 2017), and GIN (Xu et al. 2018). Compared with traditional graph embedding methods such as DeepWalk (Perozzi, Al-Rfou, and Skiena 2014) and Node2vec (Grover and Leskovec 2016), which mainly focus on graph structure, GNN can consider both node features and graph structure information. Each layer of GNN first performs message passing according to the node neighbors, and then updates the node representation through the aggregated information.

**Popular GNN-based Fraud Detection Model.** GNN-based fraud detection (GFD) can be defined as an imbalanced binary classification task that identifies outlier nodes deviating significantly from the majority of nodes. This paper targets node-level fraud detection on static graphs. For instance, GraphConsis (Liu et al. 2020) combines contextual embeddings with node features and introduces node-related relationship attention weights. CARE-GNN (Dou et al. 2020) aims at the camouflage problem in fraud detection, and designs GNN enhanced neighbor selection module. PC-GNN (Liu et al. 2021b) designs adaptive subgraph and neighbor sampling methods to fraud detection scenarios. GAGA (Wang et al. 2023) uses group aggregation to generate discriminative neighborhood information to handle low homophily problems and enhances the original feature space with learnable encoding. This brings some inspiration to our work. GTAN (Xiang et al. 2023) designs gated temporal attention networks to learn node transaction representations.

**GNN Specially Designed for Heterophily.** In fraud detection domain, nodes with different labels are often connected, which greatly limits the performance of vanilla GNNs, resulting in the emergence of a large number of studies based on heterophily.

For instance, GPRGNN (Chien et al. 2020) combines adaptive Generalized PageRank (GPR) schemes with GNNs to adapt to heterophily graphs. ACM (Luan et al. 2022) studies heterophily from the perspective of node similarity after aggregation, and then defines a new homophily index. H2-FDetector (Shi et al. 2022) identifies homophily and heterophily connections through known labeled nodes, and designs different aggregation strategies for them. BWGNN (Tang et al. 2022) designs a band-pass filter kernel function to transmit the information of each frequency band separately, thus ensuring that the spectrum energy distribution is consistent before and after message pass-

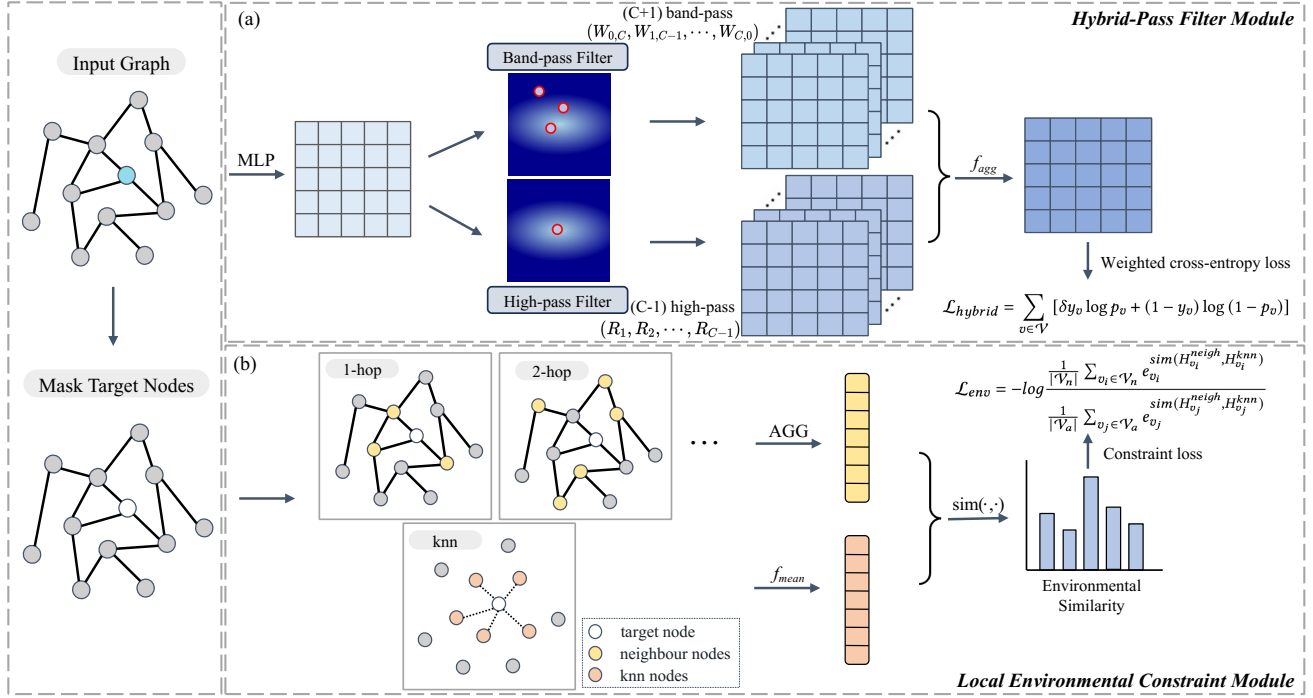


Figure 2: The above image presents an overview of our model SEC-GFD, where (a) and (b) respectively demonstrate the details of the hybrid-pass filter module and local environmental constraint module.

ing. GHRN (Gao et al. 2023a) judges and cuts heterophily edges from the perspective of graph spectrum.

## Methodology

This section begins by providing a problem definition and introducing the concepts of graph spectrum and heterophily. Subsequently, it presents the framework of the proposed model, elaborating on each module individually.

### Problem Definition

We first introduce the notations of this paper, and then give the formal problem definition. Let  $G = (\mathcal{V}, \mathcal{E}, X, Y)$  denote an attributed fraud graph, where  $\mathcal{V} = \{v_1, \dots, v_N\}$  denotes its node set,  $\mathcal{E}$  is the edge set, and  $A \in \{0, 1\}^{N \times N}$  is the adjacency matrix of  $G$  with  $A_{ij} = 1$  if there is a connection between node  $v_i$  and  $v_j$ .  $X \in R^{N \times d}$  denotes the matrix of node attributes where  $x_i \in R^d$  is the attribute vector of  $v_i$ , and  $Y$  is the corresponding labels of the nodes.

We always formulate the graph-based fraud detection (GFD) as a semi-supervised node-level binary classification problem. Most of the time, anomalies are regarded as positive with label 1, while normal nodes are viewed as negative with label 0 (Jing et al. 2021). The nodes of the whole fraud graph are divided into the labeled ones and the unlabeled ones, the former are labeled with  $Y_{train}$  and the labels of latter ones  $Y_{test}$  are invisible during training period. Thus GFD aims to learn an anomaly labeling function for our unlabeled nodes with all known information:

$$Y_{test} = f(X, A, Y_{train}). \quad (1)$$

### Graph Spectrum and Heterophily

**Graph Spectrum.** As  $A$  is the adjacency matrix of  $G$ , we let the graph laplacian matrix  $L$  be defined as  $D - A$  or as  $I - D^{-1/2}AD^{-1/2}$  (Kipf and Welling 2016), where  $I$  is an identify matrix.  $L$  is a positive semidefinite matrix and can be rewritten as  $U\Lambda U^T$ , where  $\Lambda$  is a diagonal matrix consisted of eigenvalues  $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_N$  and  $U$  is composed of the corresponding  $N$  eigenvectors  $U = (u_1, u_2, \dots, u_N)$  (Bo et al. 2021).

According to theory of graph signal processing, we can treat the eigenvectors of normalized Laplacian matrix  $U$  as the bases in Graph Fourier Transform. Suppose  $X = [x_1, x_2, \dots, x_N]$  are signals on the graph, we treat  $U^T X$  as the graph fourier transform (Hoang, Maehara, and Murata 2021) of signal  $X$ . Thus, the convolution between the signal  $x$  and convolution kernel  $g$  is as follows:

$$g * x = U((U^T g) \odot (U^T x)) = U g_w U^T x, \quad (2)$$

where  $\odot$  is the element-wise product of vectors and  $g_w$  is a diagonal matrix representing the convolutional kernel in the spectral domain. Most popular subsequent spectral GNN methods mainly focus on improving the design of  $g_w$ . And for signals  $X$ , the Rayleigh quotient represents signals' standardized variance fraction (Kim et al. 2015):

$$R(L, X) = \frac{X^T L X}{X^T X} = \frac{\sum_{(i,j) \in \mathcal{E}} (x_i - x_j)^2}{2 \sum_{i \in \mathcal{V}} x_i^2}. \quad (3)$$

Moreover, as discussed in BWGNN (Tang et al. 2022), the Rayleigh quotient can also represent the proportion of

high-frequency component energy in the spectrum, which matches the overhead variance fraction. The proof details can be found in Appendix A.

**Heterophily.** Given a set of labeled nodes and the connections that exist between them, an edge is called a heterophilic connection if its source node and destination node have different labels (i.e., anomaly and normal). For each node, the sum of its heterophily and homophily degree is equal to 1. The heterophily of the node and the whole graph can be defined separately as:

$$\begin{aligned} hetero(v) &= \frac{1}{|\mathcal{N}(v)|} |\{u : u \in \mathcal{N}(v), y_u \neq y_v\}|, \\ hetero(G) &= \frac{1}{|\mathcal{E}|} |\{\varepsilon : y_{src} \neq y_{dst}\}|, \end{aligned} \quad (4)$$

where  $|\mathcal{N}(v)|$  is the number of neighbours of node  $v$ , and  $|\mathcal{E}|$  is the total number of edges in the whole graph,  $y_{src}$  and  $y_{dst}$  are the labels of the source and destination node of the edge  $\varepsilon$  respectively.

### Overview of the Proposed Framework

Figure 2 illustrates the details of our framework. We propose a Spectrum-Enhanced and Environment-Constrained Graph Fraud Detector (SEC-GFD) to comprehensively integrate rich spectrum and label information into a fraud detector. In particular, SEC-GFD first decomposes the spectrum of the graph signal. Then it performs complicated message passing technique based on these frequency bands respectively to learn comprehensive representations. Then we design a local environmental constraint module to iteratively enhance the high-order associations between nodes' true labels and their surrounding neighbor environments.

### Hybrid-pass Filter Module

In the process of designing a hybrid band-pass filter, for purpose of ensuring the division of the spectrum into appropriate frequency bands, we utilize BWGNN (Tang et al. 2022) as our backbone. According to this, the beta wavelet transform kernel function can be written as:

$$\mathcal{W}_{p,q} = \mathbf{U} \beta_{p,q}^*(\mathbf{A}) \mathbf{U}^T = \frac{(\frac{L}{2})^p (I - \frac{L}{2})^q}{2B(p+1, q+1)}, \quad (5)$$

in which  $p+q=C$ . Then we decompose the spectrum into  $(C+1)$  frequency bands in the frequency domain, which is demonstrated below:

$$\mathcal{W} = (\mathcal{W}_{0,C}, \mathcal{W}_{1,C-1}, \dots, \mathcal{W}_{C,0}). \quad (6)$$

It is observed that  $\mathcal{W}_{p,q}$  is actually a  $C$  power polynomial, and  $C$  represents how many hops of neighbors' information are considered. Since the Laplacian matrix measures the differences of nodes, it can overcome the over-smoothing problem of vanilla GNN to a certain extent and can be extended to higher order neighbors. When unpacking  $\mathcal{W}_{p,q}$ ,  $\mathcal{W}_{C,0}$  denotes the  $C$  power of the Laplacian matrix  $L$ . The remained polynomials are the couplings of high-frequency information of neighbors below order  $C$  and low-frequency information of low-order neighbors. Furthermore, pure high-frequency information of low-order neighbors is lacked in

these polynomials, thus we expand the corresponding part and design a hybrid-pass filter. When adding the pure high-frequency information from order 1 to  $(C-1)$ , we use  $\varepsilon I - D^{-1/2} A D^{-1/2}$  or  $\varepsilon I - \tilde{A}$  (Bo et al. 2021) as the basis, where  $\varepsilon$  is a scaling hyperparameter bounded in  $[0, 1]$ . The low-order neighbor high-frequency filtering transformation is expressed as follows:

$$\begin{aligned} R_1 : \mathcal{F}_H^1 &= \varepsilon I - \tilde{A} = (\varepsilon - 1)I + L, \\ R_2 : \mathcal{F}_H^2 &= (\varepsilon I - \tilde{A})^2 = ((\varepsilon - 1)I + L)^2, \\ &\vdots \\ R_{C-1} : \mathcal{F}_H^{C-1} &= (\varepsilon I - \tilde{A})^{C-1} = ((\varepsilon - 1)I + L)^{C-1}. \end{aligned} \quad (7)$$

These  $(C-1)$  frequency bands identify the clean high-frequency information of the low-order neighbors  $\mathcal{R} = (R_1, R_2, \dots, R_{C-1})$ . So far, we obtain  $2C$  frequency bands, consisting of  $(C+1)$  band-pass frequency bands and  $(C-1)$  high-pass frequency bands for low-order neighbors. Thus we obtain the *Hybrid* spectral combination:

$$\begin{aligned} Hybrid &= Concat(\mathcal{W}, \mathcal{R}) \\ &= \overbrace{(\mathcal{W}_{0,C}, \mathcal{W}_{1,C-1}, \dots, \mathcal{W}_{C,0})}^{(C+1) \text{ Band-pass}} \overbrace{(\mathcal{R}_1, \dots, \mathcal{R}_{C-1})}^{(C-1) \text{ High-pass}}. \end{aligned} \quad (8)$$

So far there are a total of  $2C$  frequency bands. Then message transmission function is carried out on  $2C$  frequency bands separately, and finally the learned representations of each frequency band are aggregated:

$$\begin{aligned} H_0 &= MLP(X), \\ \mathcal{B}_i &= \mathcal{W}_{i,C-i} H_0 \quad \mathcal{H}_j = R_j H_0, \\ H &= f_{agg}(\mathcal{B}_0, \dots, \mathcal{B}_C, \mathcal{H}_1, \dots, \mathcal{H}_{C-1}), \end{aligned} \quad (9)$$

where  $MLP$  (Taud and Mas 2018) denotes a simple multi-layer perceptron and  $f_{agg}$  can be any aggregation function (Wu et al. 2023) such as summation or concatenation.  $H$  is the aggregated representation of hybrid-pass filter, then weighted cross-entropy loss (Zhong, Wang, and Miao 2019) is used for the training of hybrid-pass filter module:

$$\mathcal{L}_{hybrid} = \sum_{v \in \mathcal{V}} [\delta y_v \log p_v + (1 - y_v) \log (1 - p_v)], \quad (10)$$

where  $\delta$  is the ratio of anomaly labels ( $y_v = 1$ ) to normal labels ( $y_v = 0$ ) in the training set.

### Local Environmental Constraint Module

In order to improve the label utilization, and to fully consider the structure and context information of local environments, we add a local environmental constraint module. Specifically, we first mask the characteristics of the target node to prevent the influence of the characteristics of the node itself. Afterwards vanilla GNN is used to jointly learn the feature information of the target node's multi-hop neighbors and then aggregate them together. Obviously, since the

features of the target nodes are masked, thus let  $h_t^{(0)} = 0$ . The aggregation function is as follows:

$$h_t^{(l+1)} = \text{UPDATE} \left( h_t^{(l)}, \text{AGG} \left( \{h_v^{(l)} : v \in \mathcal{N}_t\} \right) \right), \quad (11)$$

$$H_t^{\text{neigh}} = \text{AGG}(h_t^1, h_t^2, \dots, h_t^L), \quad (12)$$

where  $H_t^{\text{neigh}}$  is the final neighbor aggregation representation. For purpose of exploring the essential association between neighbors with different hop counts, a simple SGC (Wu et al. 2019) is used without adding any nonlinear functions. Moreover, in order to alleviate the over-smoothing effect (Chen et al. 2020),  $L$  is generally selected as 2 or 3 in the experiments.

Then we progress graph augmentation and construct KNN view of fraud graph according to the features. Concretely, the  $k$ -nearest neighbors (Dong, Moses, and Li 2011) of each node are found according to the feature cosine similarity. To ensure the rationality of local environmental features, we do not add any nonlinear functions, but employ the average pooling operation.

$$H_t^{\text{knn}} = f_{\text{mean}}(\{\mathbf{x}_u | \forall u \in K_t\}) = \frac{1}{|K_t|} \sum_{u \in K_t} \mathbf{x}_u. \quad (13)$$

So far, we have obtained the multi-hop neighbor information of the target nodes and their  $k$ -nearest neighbor environmental feature information. In order to realize the local environmental constraints, this paper judges based on common senses that, the multi-hop neighbor features learned by normal nodes and the local environmental features should be as similar as possible, while the two features learned by anomaly nodes should have a large deviation. Following the formula of InfoNCE loss (You et al. 2020) in Contrastive Learning, we design a constraint loss function  $\mathcal{L}_{\text{env}}$  specially adapted to the above judgments:

$$\mathcal{L}_{\text{env}} = -\log \frac{\frac{1}{|\mathcal{V}_n|} \sum_{v_i \in \mathcal{V}_n} \frac{\text{sim}(H_{v_i}^{\text{neigh}}, H_{v_i}^{\text{knn}})}{e_{v_i}}}{\frac{1}{|\mathcal{V}_a|} \sum_{v_j \in \mathcal{V}_a} \frac{\text{sim}(H_{v_j}^{\text{neigh}}, H_{v_j}^{\text{knn}})}{e_{v_j}}}, \quad (14)$$

where  $\mathcal{V}_n$  and  $\mathcal{V}_a$  represent the set of normal nodes and anomaly nodes in the training set respectively, and  $\text{sim}(\cdot, \cdot)$  is the simple cosine similarity operator.

## Fraud Detection

For the sake of jointly learning hybrid frequency-band representations and applying local environmental constraints, we carefully design the objective function of SEC-GFD. Specifically, we construct a weighted binary cross-entropy loss function and an InfoNCE-like contrastive loss to realize the constraints of the surrounding environments on the target nodes. The final loss function is as follows:

$$\mathcal{L} = \alpha \mathcal{L}_{\text{hybrid}} + (1 - \alpha) \mathcal{L}_{\text{env}}. \quad (15)$$

In order to better combine the two loss functions, a hyperparameter  $\alpha \in [0, 1]$  is added to balance their influences.

Dataset	Nodes	Edges	Anomaly(%)	Features
Amazon	11,944	4,398,392	6.87	25
YelpChi	45,954	3,846,979	14.53	32
T-Finance	39,357	21,222,543	4.58	10
T-Social	5,781,065	73,105,508	3.01	10

Table 1: Statistics of four datasets.

## Experiments

In this section, we evaluate the performance of our proposed approach on real-world datasets through a series of experiments, and compare the results with those of state-of-the-art baselines to demonstrate the effectiveness of our approach. In addition, we conduct experiments on the response of heterophily edges to distinct filters, as well as experiments on ablation and hyperparameter analysis.

### Experimental Settings

**Dataset.** We conduct experiments on four real-world datasets targeted at fraud detection scenarios. Overall, they are Amazon (McAuley and Leskovec 2013), YelpChi (Rayana and Akoglu 2015), and two recently published transaction datasets, i.e., T-Finance and T-Social (Tang et al. 2022). The YelpChi dataset consists of filtered and recommended hotel and restaurant reviews from Yelp. The Amazon dataset includes product reviews in the music instrument category. The T-Finance dataset and T-Social dataset detect anomalous accounts in transaction networks and social networks respectively, and they share the same 10-dimensional features. The T-Social dataset is a large-scale dataset, several tens to hundreds of times larger in scale compared to the other datasets. The statistics of these datasets are summarized in Table 1.

**BaseLines for Comparison.** We conduct comparative experiments between SEC-GFD and advanced baselines that belong to three groups. The first group consists of methods based on homophily GNN models, including GCN (Kipf and Welling 2016), GraphSAGE (Hamilton, Ying, and Leskovec 2017), GAT (Velivcković et al. 2017) and GIN (Xu et al. 2018). The second group comprises state-of-the-art methods for graph-based fraud detection, namely GraphConsis (Liu et al. 2020), CARE-GNN (Dou et al. 2020), PC-GNN (Liu et al. 2021b) and GAGA (Wang et al. 2023). The third group contains GNN-based fraud detection models specifically designed for heterophily, including ACM (Luan et al. 2022), H2-FDetector (Shi et al. 2022), BWGNN (Tang et al. 2022) and GDN (Gao et al. 2023b). More details on description and implementation of the benchmarks are provided in Appendix B.

**Metrics.** As graph anomaly detection poses a class-imbalanced classification problem, this paper utilizes two widely adopted metrics: F1-macro and AUC. F1-macro considers the weighted average of F1 scores across multiple classes, and AUC is the area under the ROC Curve.

**Settings.** In our proposed method SEC-GFD, for the Amazon, YelpChi, and T-Finance datasets, the hidden layer dimension is set to 64, and the high-frequency signal neighbor order  $C$  is set to 2. For the T-Social dataset, the hidden layer

Methods		Datasets							
		YelpChi		Amazon		T-Finance		T-Social	
		F1-macro	AUC	F1-macro	AUC	F1-macro	AUC	F1-macro	AUC
Homophily GNN	GCN	53.21	56.80	63.52	80.14	71.47	66.31	61.73	87.84
	GraphSAGE	64.34	73.61	75.25	86.73	55.27	69.52	60.49	72.65
	GAT	55.86	58.92	82.41	89.73	54.27	75.29	73.72	88.61
	GIN	63.42	75.28	71.14	81.38	65.72	81.63	63.38	81.33
GFD Algorithm	GraphConsis	57.91	69.55	78.46	87.27	73.58	91.42	58.32	73.89
	CARE-GNN	63.58	79.42	83.31	91.57	65.37	91.93	55.92	68.78
	PC-GNN	64.75	77.14	91.73	95.63	77.25	92.83	57.21	71.27
	GAGA	<u>76.71</u>	89.77	90.31	95.61	85.13	94.72	76.64	87.93
Heterophily GNN	ACM	69.72	88.28	81.83	93.69	85.48	<u>96.02</u>	81.27	90.79
	H2-FDetector	69.38	88.63	83.84	96.41	<u>87.39</u>	95.41	78.89	88.56
	BWGNN	76.66	90.56	<u>91.68</u>	<u>97.52</u>	84.59	95.29	84.58	95.27
	GDN	76.05	<u>90.79</u>	90.68	97.09	86.12	94.28	80.63	89.35
ours	SEC-GFD	<b>77.73</b>	<b>91.39</b>	<b>92.35</b>	<b>98.23</b>	<b>89.86</b>	<b>96.32</b>	<b>87.74</b>	<b>96.11</b>

Table 2: Fraud detection performance on four real-world datasets in terms of F1-macro and AUC value. (The bold and italicized values in the table represent the best and second-best performance, respectively.)

dimension is set to 16, the order C of high-frequency signal neighbors is set to 5. The experimental results show that, for such a large-scale graph, fusing high-frequency information of higher-order neighbors can learn better features. In the experiments, to ensure fairness, the size of training/validation/testing set of the datasets is set to 0.4/0.2/0.4 for all the compared methods. Our method is implemented with the DGL library in Pytorch, and for other methods, public source codes are utilized for implementation. Furthermore, we run for 100 epochs on the four datasets for all methods.

## Experimental Results

We conduct a comprehensive comparison of our approach with vanilla homophily GNN models, state-of-the-art graph-based fraud detection models and novel GNNs targeted at heterophily. The corresponding results are shown in Table 2.

First of all, It is observed that SEC-GNN outperforms all vanilla homophily GNN models, *i.e.*, GCN, GraphSAGE, GAT and GIN. For instance, when compared with GIN, SEC-GFD achieves improvements of 13.37% and 26.41% in terms of F1-macro and AUC on Amazon. To delve reasons, these models can be regarded as low-pass filters in essence and only focus on low frequency information. While our SEC-GNN integrates rich information of distinct frequency bands adapting to heterophily.

When compared with state-of-the-art graph-based fraud detection methods, these methods still have a large gap with our approach. Among them, GraphConsis demonstrates the worst comprehensive performance. It combines the characteristics of nodes with the surrounding environments to learn relationship attention weights, and it is essentially a low-pass filter. CARE-GNN and PC-GNN achieve some improvements in the effectiveness. This is because they are aware of heterophily edges and camouflage problems, and design corresponding neighbor selection module. GAGA introduces a group aggregation module to generate distinguishable multi-hop neighborhood information. Although it demonstrates substantial overall improvements, it still ex-

hibits a significant disparity when compared to our proposed method. Specically, when compared with GAGA, SEC-GFD achieves gains of 2.04% and 2.62% in terms of F1-macro and AUC on YelpChi. The reason behind this can be attributed to that our method additionally learns complicated high-frequency information in addition to learning multi-hop neighbor information.

Finally, when compared with novel heterophily GNNs, experimental results in Table 2 show that our proposed method consistently outperforms them on all metrics across the four datasets. To be specific, SEC-GFD attains improvements of 2.47% and 1.71% in terms of F1-macro and AUC on T-Finance. As to T-Social, SEC-GFD achieves gains of 3.78% and 0.84% in terms of F1-macro and AUC. Among these models, ACM studies heterophily in terms of feature similarity, H2-FDetector identifies homophily and heterophily connections and applies different aggregation strategies, BWGNN designs a band-pass filter to transmit information in multiple frequency bands, and GDN divides the node features into anomaly pattern and neighborhood pattern. The superior performance of SEC-GFD indicates the effectiveness of the proposed hybrid filtering module and local environmental constraint module.

## Heterophily Edges Clipping Analysis

To explore the effect of heterophily edge pruning on GNNs with different filter properties, we conduct extensive experiments. In this paper, three message passing mechanisms with low-pass, high-pass, and band-pass filtering characteristics are designed for experiments. The low-pass and high-pass filters use the adjacency matrix and normalized Laplacian matrix respectively for message passing, and the band-pass filter uses the beta wavelet filter function from BWGNN, with two layers of neighbor information are considered. The experimental results on Amazon and Yelpchi are illustrated in Figure 3.

Considering the difficulty in discriminating heterophily edges, in many cases only some heterophily edges con-



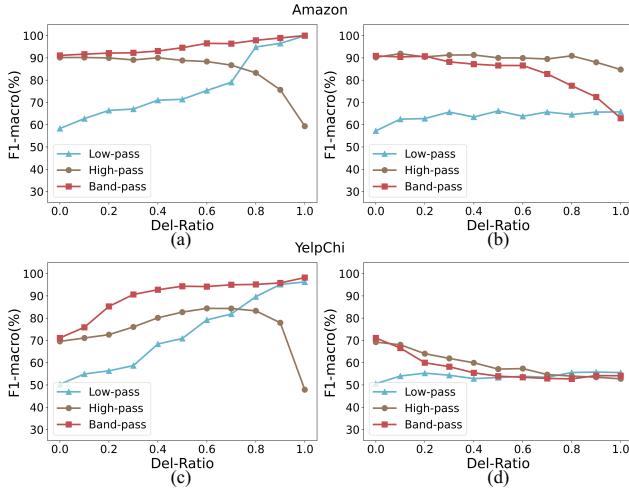


Figure 3: Filters’ performance with different heterophily edges’ deletion ratio on Amazon and YelpChi. Fig (a) and Fig (c) denote deletion of heterophily edges on the entire graph, while Fig (b) and Fig (d) denote deletion of heterophily edges appearing in training graph.

nected with labeled nodes on the training set can be identified. Thus two sets of experiments are conducted in this paper. As observed in Figure 3, when clipping the heterophily edges of the whole graph, the performance of low-pass and band-pass filtering improves with the ratio of heterophily edge clipping. While the performance of high-pass filtering shows a monotonous decrease. And the performance of the band-pass filter keeps the best. Moreover, when clipping heterophily edges appearing in the training set, results are completely different. Specifically, low-pass filter maintains a poor performance, while band-pass and high-pass filter exhibit different degrees of declining trends. This suggests that indiscriminate removal of heterophily edges is not necessarily beneficial.

### Effectiveness of Each Component of SEC-GFD

To explore how each component contributes to the performance of SEC-GFD, we systematically remove certain parts of the framework for ablation study. The results are illustrated in Table 3.

Firstly, in order to explore the performance of the two parts of the hybrid filter, we omit band-pass bands and low-order pure high-pass bands respectively. When band-pass bands are excluded, the performance drops by 2.06% and 2.67% in terms of AUC and F1-macro. And the performance decreases 1.13% and 1.27% in terms of AUC and F1-macro when low-order high-pass bands are removed. It is obviously that band-pass bands are more important, this is because band-pass bands consider high frequency and lower frequency information together. Afterwards, we delete the local environmental constraint, then performance drops by 0.86% and 1.31% in terms of AUC and F1-macro. However, it still outperforms other state-of-the-art methods, this also reflects the effectiveness of our hybrid filter. Thus the hybrid filter can be equipped as a general backbone to various other

	Amazon		T-Finance	
	F1-macro	AUC	F1-macro	AUC
SEC-GFD	<b>92.27</b>	<b>98.19</b>	<b>89.74</b>	<b>96.34</b>
w/o high-pass	91.48	97.28	88.65	95.05
w/o band-pass	90.29	96.31	86.79	93.77
w/o env	<u>91.88</u>	97.14	88.14	<u>95.21</u>

Table 3: Ablation Performance.(The bold and italicized values denote the best and runner-up performance.)




Figure 4: Performance with different order C.

GNN-based fraud detectors to improve performance.

### Selection of C for high-order neighbors

In the previous discussion, based on the formula of beta wavelet, we judge that C represents the order of neighbor information considered in the network training process. Therefore, the order C is an extremely important parameter. Figure 4 shows the changes of F1 and AUC values with the changes of C from 1 to 6 on four real-world datasets.

It is observed that on the three smaller datasets, the improvement of the effect is no longer obvious when C exceeds two. While on large-scale dataset T-Social, which is nearly a hundred times larger than other datasets, there is still a significant performance improvement until C reaches five. A possible reason is that, larger fraud graph may contain more complicated inter-node interactions and more common spread of fraudulent behavior across multi-hop neighbors.

## Conclusion

This paper reviews the research on highly heterophily graphs, especially in fraud detection scenarios, and further explores the association between heterophily edges and spectral domain energy distribution. The experimental results prove that it is not necessary to discriminate and cut heterophily edges, but to make full use of each frequency band for information transmission to achieve excellent results. To this end, we have designed a spectrum-enhanced and environment-constrained graph-based fraud detector SEC-GFD. The main body is a hybrid filtering module, which contains multiple frequency bands with multi-order neighbors. Then it is supplemented with a local environmental constraint module to improve the utilization of node labels and the adaptive relationships between nodes and their surrounding environments.

## Acknowledgments

This work is supported in part by the Fundamental Research Funds for the Central Universities under Grant 2022RC026; in part by the CCF-NSFOCUS Open Fund; in part by the NSFC Program under Grant 62202042, Grant U20A6003, Grant 62076146, Grant 62021002, Grant U19A2062, Grant 62127803, Grant U1911401 and Grant 6212780016; in part by Industrial Technology Infrastructure Public Service Platform Project ‘Public Service Platform for Urban Rail Transit Equipment Signal System Testing and Safety Evaluation’ (No. 2022-233-225), Ministry of Industry and Information Technology of China.

## References

- Akoglu, L.; Tong, H.; and Koutra, D. 2015. Graph based anomaly detection and description: a survey. *Data mining and knowledge discovery*, 29: 626–688.
- Bo, D.; Wang, X.; Shi, C.; and Shen, H. 2021. Beyond low-frequency information in graph convolutional networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 3950–3957.
- Chen, D.; Lin, Y.; Li, W.; Li, P.; Zhou, J.; and Sun, X. 2020. Measuring and relieving the over-smoothing problem for graph neural networks from the topological view. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, 3438–3445.
- Chien, E.; Peng, J.; Li, P.; and Milenkovic, O. 2020. Adaptive universal generalized pagerank graph neural network. *arXiv preprint arXiv:2006.07988*.
- Dong, W.; Moses, C.; and Li, K. 2011. Efficient k-nearest neighbor graph construction for generic similarity measures. In *Proceedings of the 20th international conference on World wide web*, 577–586.
- Dou, Y.; Liu, Z.; Sun, L.; Deng, Y.; Peng, H.; and Yu, P. S. 2020. Enhancing graph neural network-based fraud detectors against camouflaged fraudsters. In *Proceedings of the 29th ACM international conference on information & knowledge management*, 315–324.
- Gao, Y.; Wang, X.; He, X.; Liu, Z.; Feng, H.; and Zhang, Y. 2023a. Addressing heterophily in graph anomaly detection: A perspective of graph spectrum. In *Proceedings of the ACM Web Conference 2023*, 1528–1538.
- Gao, Y.; Wang, X.; He, X.; Liu, Z.; Feng, H.; and Zhang, Y. 2023b. Alleviating structural distribution shift in graph anomaly detection. In *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining*, 357–365.
- Grover, A.; and Leskovec, J. 2016. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, 855–864.
- Hamilton, W.; Ying, Z.; and Leskovec, J. 2017. Inductive representation learning on large graphs. *Advances in neural information processing systems*, 30.
- Hoang, N.; Maehara, T.; and Murata, T. 2021. Revisiting graph neural networks: Graph filtering perspective. In *2020 25th International Conference on Pattern Recognition (ICPR)*, 8376–8383. IEEE.
- Hooi, B.; Shin, K.; Song, H. A.; Beutel, A.; Shah, N.; and Faloutsos, C. 2017. Graph-based fraud detection in the face of camouflage. *TKDD*, 11(4): 1–26.
- Jin, M.; Liu, Y.; Zheng, Y.; Chi, L.; Li, Y.-F.; and Pan, S. 2021. Anemone: Graph anomaly detection with multi-scale contrastive learning. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, 3122–3126.
- Jing, R.; Tian, H.; Zhou, G.; Zhang, X.; Zheng, X.; and Zeng, D. D. 2021. A GNN-based Few-shot learning model on the Credit Card Fraud detection. In *2021 IEEE 1st International Conference on Digital Twins and Parallel Intelligence (DTPI)*, 320–323. IEEE.
- Kim, T.; Wang, Y.; Sahinoglu, Z.; Wada, T.; Hara, S.; and Qiao, W. 2015. A Rayleigh quotient-based recursive total-least-squares online maximum capacity estimation for lithium-ion batteries. *IEEE Transactions on Energy Conversion*, 30(3): 842–851.
- Kipf, T. N.; and Welling, M. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*.
- Liu, Y.; Ao, X.; Qin, Z.; Chi, J.; Feng, J.; Yang, H.; and He, Q. 2021a. Pick and choose: a GNN-based imbalanced learning approach for fraud detection. In *WWW*, 3168–3177.
- Liu, Y.; Ao, X.; Qin, Z.; Chi, J.; Feng, J.; Yang, H.; and He, Q. 2021b. Pick and choose: a GNN-based imbalanced learning approach for fraud detection. In *Proceedings of the web conference 2021*, 3168–3177.
- Liu, Z.; Dou, Y.; Yu, P. S.; Deng, Y.; and Peng, H. 2020. Alleviating the inconsistency problem of applying graph neural network to fraud detection. In *Proceedings of the 43rd international ACM SIGIR conference on research and development in information retrieval*, 1569–1572.
- Luan, S.; Hua, C.; Lu, Q.; Zhu, J.; Zhao, M.; Zhang, S.; Chang, X.-W.; and Precup, D. 2022. Revisiting heterophily for graph neural networks. *Advances in neural information processing systems*, 35: 1362–1375.
- McAuley, J. J.; and Leskovec, J. 2013. From amateurs to connoisseurs: modeling the evolution of user expertise through online reviews. In *Proceedings of the 22nd international conference on World Wide Web*, 897–908.
- Perozzi, B.; Al-Rfou, R.; and Skiena, S. 2014. Deepwalk: Online learning of social representations. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, 701–710.
- Pourhabibi, T.; Ong, K.-L.; Kam, B. H.; and Boo, Y. L. 2020. Fraud detection: A systematic literature review of graph-based anomaly detection approaches. *Decision Support Systems*, 133: 113303.
- Qin, Z.; Liu, Y.; He, Q.; and Ao, X. 2022. Explainable Graph-based Fraud Detection via Neural Meta-graph Search. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, 4414–4418.



- Rayana, S.; and Akoglu, L. 2015. Collective opinion spam detection: Bridging review networks and metadata. In *Proceedings of the 21th acm sigkdd international conference on knowledge discovery and data mining*, 985–994.
- Shi, F.; Cao, Y.; Shang, Y.; Zhou, Y.; Zhou, C.; and Wu, J. 2022. H2-fdetector: A gnn-based fraud detector with homophilic and heterophilic connections. In *Proceedings of the ACM Web Conference 2022*, 1486–1494.
- Tang, J.; Li, J.; Gao, Z.; and Li, J. 2022. Rethinking graph neural networks for anomaly detection. In *International Conference on Machine Learning*, 21076–21089. PMLR.
- Taud, H.; and Mas, J. 2018. Multilayer perceptron (MLP). *Geomatic approaches for modeling land change scenarios*, 451–455.
- Velivcković, P.; Cucurull, G.; Casanova, A.; Romero, A.; Lio, P.; and Bengio, Y. 2017. Graph attention networks. *arXiv preprint arXiv:1710.10903*.
- Wang, H.; Zhou, C.; Wu, J.; Dang, W.; Zhu, X.; and Wang, J. 2018. Deep structure learning for fraud detection. In *ICDM*, 567–576.
- Wang, Y.; Zhang, J.; Huang, Z.; Li, W.; Feng, S.; Ma, Z.; Sun, Y.; Yu, D.; Dong, F.; Jin, J.; et al. 2023. Label Information Enhanced Fraud Detection against Low Homophily in Graphs. In *Proceedings of the ACM Web Conference 2023*, 406–416.
- Wu, F.; Souza, A.; Zhang, T.; Fifty, C.; Yu, T.; and Weinberger, K. 2019. Simplifying graph convolutional networks. In *International conference on machine learning*, 6861–6871. PMLR.
- Wu, H.; Xion, W.; Xu, F.; Luo, X.; Chen, C.; Hua, X.-S.; and Wang, H. 2023. PastNet: Introducing Physical Inductive Biases for Spatio-temporal Video Prediction. *arXiv preprint arXiv:2305.11421*.
- Xiang, S.; Zhu, M.; Cheng, D.; Li, E.; Zhao, R.; Ouyang, Y.; Chen, L.; and Zheng, Y. 2023. Semi-supervised credit card fraud detection via attribute-driven graph representation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 14557–14565.
- Xu, F.; Wang, N.; Wen, X.; Gao, M.; Guo, C.; and Zhao, X. 2023. Few-shot Message-Enhanced Contrastive Learning for Graph Anomaly Detection. *arXiv preprint arXiv:2311.10370*.
- Xu, K.; Hu, W.; Leskovec, J.; and Jegelka, S. 2018. How powerful are graph neural networks? *arXiv preprint arXiv:1810.00826*.
- You, Y.; Chen, T.; Sui, Y.; Chen, T.; Wang, Z.; and Shen, Y. 2020. Graph contrastive learning with augmentations. *Advances in neural information processing systems*, 33: 5812–5823.
- Zhang, G.; Wu, J.; Yang, J.; Beheshti, A.; Xue, S.; Zhou, C.; and Sheng, Q. Z. 2021. Fraudre: Fraud detection dual-resistant to graph inconsistency and imbalance. In *ICDM*, 867–876.
- Zhang, Z.; Wan, J.; Zhou, M.; Lai, Z.; Tessone, C. J.; Chen, G.; and Liao, H. 2023. Temporal burstiness and collaborative camouflage aware fraud detection. *Information Processing & Management*, 60(2): 103170.
- Zhong, P.; Wang, D.; and Miao, C. 2019. An affect-rich neural conversational model with biased attention and weighted cross-entropy loss. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 7492–7500.
- Zhu, J.; Rossi, R. A.; Rao, A.; Mai, T.; Lipka, N.; Ahmed, N. K.; and Koutra, D. 2021. Graph neural networks with heterophily. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, 11168–11176.