

Chi-Square Wavelet Graph Neural Networks for Heterogeneous Graph Anomaly Detection

Xiping Li*
The Chinese University of
Hong Kong
Hong Kong SAR, China
lihsiping@gmail.com

Xiangyu Dong
The Chinese University of
Hong Kong
Hong Kong SAR, China
xydong@se.cuhk.edu.hk

Xingyi Zhang
Mohamed bin Zayed
University of Artificial
Intelligence
Abu Dhabi, United Arab
Emirates
xingyi.zhang@mbzuai.ac.ae

Kun Xie
The Chinese University of
Hong Kong
Hong Kong SAR, China
kxie@se.cuhk.edu.hk

Yuanhao Feng
Tencent, WeChat Pay
Shenzhen, China
danielfeng@tencent.com

Bo Wang
Tencent, WeChat Pay
Shenzhen, China
pollowang@tencent.com

Guilin Li
Tencent, WeChat Pay
Shenzhen, China
guilinli@tencent.com

Wuxiong Zeng
Tencent, WeChat Pay
Shenzhen, China
philzeng@tencent.com

Xiujun Shu
Tencent, WeChat Pay
Shenzhen, China
xiujunshu@tencent.com

Sibo Wang[†]
The Chinese University of
Hong Kong
Hong Kong SAR, China
swang@se.cuhk.edu.hk

Abstract

Graph Anomaly Detection (GAD) in heterogeneous networks presents unique challenges due to node and edge heterogeneity. Existing Graph Neural Network (GNN) methods primarily focus on homogeneous GAD and thus fail to address three key issues: (C1) Capturing abnormal signal and rich semantics across diverse meta-paths; (C2) Retaining high-frequency content in HIN dimension alignment; and (C3) Learning effectively from difficult anomaly samples with class imbalance. To overcome these, we propose ChiGAD, a spectral GNN framework based on a novel Chi-Square filter, inspired by the wavelet effectiveness in diverse domains. Specifically, ChiGAD consists of: (1) Multi-Graph Chi-Square Filter, which captures anomalous information via applying dedicated Chi-Square filters to each meta-path graph; (2) Interactive Meta-Graph Convolution, which aligns features while preserving high-frequency information and incorporates heterogeneous messages by a unified Chi-Square Filter; and (3) Contribution-Informed Cross-Entropy Loss, which prioritizes difficult anomalies to address class imbalance. Extensive experiments on public and industrial datasets show that ChiGAD outperforms state-of-the-art models on multiple metrics. Additionally, its homogeneous variant, ChiGNN, excels on seven GAD datasets, validating the effectiveness of Chi-Square filters. Our code is available at <https://github.com/HsipingLi/ChiGAD>.

*This work was done during an internship at CUHK and Tencent.

[†]Corresponding author



This work is licensed under a Creative Commons Attribution 4.0 International License.
KDD '25, Toronto, ON, Canada

© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1454-2/2025/08
<https://doi.org/10.1145/3711896.3736877>

CCS Concepts

• **Computing methodologies** → **Anomaly detection.**

Keywords

Heterogeneous Graph Anomaly Detection, Chi-Square Wavelet, Spectral Graph Neural Networks

ACM Reference Format:

Xiping Li, Xiangyu Dong, Xingyi Zhang, Kun Xie, Yuanhao Feng, Bo Wang, Guilin Li, Wuxiong Zeng, Xiujun Shu, and Sibow Wang. 2025. Chi-Square Wavelet Graph Neural Networks for Heterogeneous Graph Anomaly Detection. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.2 (KDD '25)*, August 3–7, 2025, Toronto, ON, Canada. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3711896.3736877>

1 Introduction

Graph Anomaly Detection (GAD) [33] has emerged as a critical research focus of graph learning, with vital applications ranging from identifying risky merchants [25] to detecting fake news, [32, 45], and malicious users [27, 37]. Graph Neural Networks (GNNs) [21, 36] have become the preferred solution for GAD tasks [14, 34], given their inherent ability to process graph-structured data. Since node-level GAD is fundamentally a binary node classification task, thorough assessments of various traditional GNN models [15, 21] have been conducted. Recent spectral GNNs [5, 14, 34], especially those employing wavelet filters [2, 28, 34], have shown great superiority in homogeneous GAD tasks. Yet, these approaches invariably assume unified node and edge types of fraud graphs, ignoring the scenarios characterized by Heterogeneous Information Networks (HINs). Despite the successful application of GNN-based models in various HINs, including social network dynamics [22, 44], link prediction [7, 23], and recommendations [4, 12, 24, 31, 42],

anomaly detection in heterogeneous networks remains understudied. Considering the significance of GAD in critical HINs such as financial networks, a specialized framework for heterogeneous GAD is urgently needed.

Learning on HINs is inherently challenging due to both the node and edge heterogeneity [43]. Recently, a series of meta-path-based GNNs [13, 20, 38, 43, 46] and multi-graph filters [3, 17] have been proposed to address the issues arising from heterogeneity. However, it is worth noting that they overlook the following critical challenges due to not being specifically designed for GAD tasks. Particularly, **(C1) Capturing abnormal signals and rich semantics across diverse meta-paths:** In GAD applications such as fraud detection, meta-paths exhibit distinct frequency characteristics crucial for anomaly identification. For instance, in a financial network, transaction-based meta-paths may show irregular bursts, while social interactions remain steady. A generic filter cannot capture such differences across frequency bands. Previous works [3, 17] adopt a unified signal processing framework, thereby overlooking the nuanced frequency energy distributions inherent in different meta-path graphs. **(C2) Retaining high-frequency content in HIN dimension alignment:** The retention of high-frequency content is essential to avert over-smoothing and performance degradation in homogeneous GAD tasks [14, 34]; additionally, the node heterogeneity inherent in HIN engenders dimensional disparity, compelling nodes to depend on a prior dimensional alignment technique for feasible representation learning within heterogeneous graphs [43]. Nevertheless, the risky potential for the alignment method to induce loss of high-frequency content within these signals remains an unexplored domain since they directly process graph signals. **(C3) Learning effectively from difficult anomaly samples with class imbalance:** In GAD tasks, the inherent class imbalance poses significant challenges. This imbalance makes it difficult for models to learn subtle anomaly characteristics. While existing studies [11] have acknowledged this challenge, a comprehensive strategy addressing class imbalance while effectively learning from rare and difficult anomaly samples is still lacking.

Inspired by the demonstrated effectiveness of wavelet in various domains [29, 40], especially in homogeneous GAD [2, 28, 34], we first employ the Chi-Square distribution to design a novel graph filter, namely, **Chi-Square filter**. This novel graph wavelet filter not only inherits the proven wavelet properties crucial for GAD tasks [2, 28, 34], but also possesses a unique property of additivity compared to other wavelets, which makes Chi-Square naturally suitable for collaborative graph representation learning across diverse meta-paths within HIN, as elaborated in Section 4.1.

Furthermore, motivated by the aforementioned three crucial challenges, we introduce a novel framework for heterogeneous graph anomaly detection, namely **ChiGAD**, which takes our proposed Chi-Square filter as a basic component and is comprised of three essential components: **(1) Multi-Graph Chi-Square Filter:** It processes different node types independently. For each type of nodes, it leverages the additivity of Chi-Square to collaboratively capture fine-grained semantics across diverse meta-paths by assigning the most aligned Chi-Square filters according to the spectral energy distributions of representative meta-path graphs; **(2) Interactive Meta-Graph Convolution:** It addresses representation dimension disparities in heterogeneous graphs by using a linear

layer to align feature dimensions, justified by its ability to preserve high-frequency components to resist over-smoothing, which is helpful to heterogeneous GAD tasks, as proven in Section 4.3 in theory and analyzed in Appendix 5.4. It aligns node representations across types and employs a unified Chi-Square filter on a transformed homogeneous graph to learn final nodal representations, integrating diverse frequency bands; **(3) Contribution-Informed Cross-Entropy Loss:** It assigns higher weights to anomalies based on their contributions, ensuring the model prioritizes difficult samples while maintaining focus on the scarcity of anomalous data.

In summary, our key contributions are as follows:

- To the best of our knowledge, we are the first to use Chi-Square wavelets for spectral GNNs. In addition to properties possessed by general wavelets, the unique additivity of Chi-Square wavelet enables collaborative learning across multiple meta-paths, making them seamlessly suited for heterogeneous GAD.
- We propose ChiGAD, a novel spectral GNN framework for heterogeneous GAD, which captures rich semantics across diverse meta-paths and broad frequency bands, retains the high-frequency content in alignment, and effectively learns from difficult anomalies while mitigating the class-imbalance problem.
- We provide theoretical analysis and empirical validation of high-frequency retention during dimensional alignment, which is crucial for mitigating over-smoothing in GAD tasks. Both theoretical findings in Theorem 1 and experimental results shown in Appendix 5.4 indicate that linear layer is a good choice for dimensional alignment in heterogeneous GAD.
- Experimental results on both industrial and public heterogeneous datasets demonstrate the effectiveness of ChiGAD in heterogeneous GAD tasks, showing significant improvements over state-of-the-art baselines across diverse metrics, including AUROC, AUPRC, F1-macro, and Recall.
- We develop ChiGNN, a homogeneous variant of ChiGAD, which achieves superior performance on seven public datasets, further validating the effectiveness of the Chi-Square filter.

2 Related Work

2.1 GNN-based Graph Anomaly Detection

GNN-based models are popular for node-level GAD in homogeneous graphs [9, 10, 33]. However, their effectiveness is limited by the low-pass filtering of Graph Convolutional Networks [14, 21, 34] and the heterophily phenomenon [14, 39], where anomalies connect to normal nodes. Several GNN-based anomaly detectors [5, 14, 34, 39] have been proposed to address these challenges. Specifically, BWGNN [34] incorporates Beta wavelet band-pass filters to detect graph anomalies, leveraging the right-shift phenomenon in spectral energy. XGBGraph [33] aggregates neighborhood features within an XGBoost model to capture both structural and feature cues. GHRN [14] mitigates heterophily by pruning edges identified via a high-frequency indicator, while GAGA [39] improves fraud detection in low-homophily graphs through group aggregation and learnable encoding. ConsisGAD [6] addresses limited supervision with pseudo-label generation. Yet, these models invariably overlook heterogeneous scenarios that better reflect real-world conditions.

2.2 Heterogeneous Graph Learning.

Diverse investigations [3, 13, 17, 38, 43, 46] have been introduced to enable graph representation learning in heterogeneous networks, addressing the limitations of traditional GNN-based models. Meta-path-based models provide foundational solutions. Specifically, HAN [38] uses a hierarchical attention mechanism combined with multiple meta-paths to capture node-level and semantic-level importance in HINs. MAGNN [13] enhances this approach by encoding all meta-path information rather than just endpoints, addressing the limitation of ignoring intermediate nodes. HetGNN [43] leverages random walks with restart and Bi-LSTM [30] to aggregate node features across node types. Recently, spectral GNN extensions to HINs have emerged. MGNN [3] proposes a polynomial multi-graph filter where each monomial term serves as a non-commutative multiplication of graph shifting operators. PSHGCN [17] introduces a positive semi-definite constraint to enhance model performance. Yet, these approaches are not specifically designed for GAD tasks.

3 Preliminaries

3.1 Heterogeneous Information Networks

A heterogeneous information network (HIN) is defined as a graph $G_M = (\mathcal{V}, \mathcal{A}, \mathcal{X}, O_V, R_E)$, where $\mathcal{V} = \bigcup_{o=1}^{|O_V|} \{V_o\}$ represents multiple types of nodes, each enriched with attributes $\mathcal{X} = \{X_o\}_{o=1}^{|O_V|}$; $\mathcal{A} = \{A_r\}_{r=1}^{|R_E|}$ includes the adjacency matrices corresponding to diverse relations, where $(A_r)_{ij} = 1$ if there exists an edge between nodes $i, j \in \mathcal{V}$ under relation r , otherwise $(A_r)_{ij} = 0$; O_V and R_E respectively denote the sets of node types and relation types, and $|O_V| + |R_E| > 2$. For HINs, it is usually observed that there is a dimensional disparity in node attributes $X_o \in \mathbb{R}^{|V_o| \times d_o}$, i.e. $d_i \neq d_j$ for $i \neq j$, $i, j \in O_V$; in addition, $A_r \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{V}|}$ for $r = 1, 2, \dots, |R_E|$.

3.2 Meta-Path Graphs

The meta-path is a path pattern defined on a given heterogeneous information network G_M , describing an alternate sequence of node types and edge types. A meta-path P of length l can be defined as:

$$P = o_1 \xrightarrow{r_1} o_2 \xrightarrow{r_2} \dots \xrightarrow{r_l} o_{l+1}, \quad (1)$$

where $\{o_i\}_{i=1}^{l+1}$ and $\{r_i\}_{i=1}^l$ are collections of node types and relations, respectively. In addition, a meta-path graph G_P is a subgraph extracted from the meta-graph G_M based on P defined in Equation 1, containing nodes and edges that conform to P .

The meta-path reveals the high-level semantic information hidden in HINs. For instance, in an academic network, the meta-path:

$$\text{author} \xrightarrow{\text{compose}} \text{paper} \xrightarrow{\text{composed by}} \text{author}$$

captures co-authoring semantics among authors. The diversity of meta-path composition is an inherent property of HINs, and effectively mining rich information across meta-paths way poses challenges for HIN-related tasks [38].

3.3 Dynamics of Information Diffusion

Meta-paths capture complex relations between nodes in HINs. Given an HIN G_M , each meta-path P as defined in Equation 1 varies based on both the composition and length of its sequence

r_1, r_2, \dots, r_l . Following prior work [3], information diffusion along a meta-path P refers to an alternate execution of message propagation and aggregation, formally expressed as:

$$\mathbf{X}' = S_{r_l} \cdots S_{r_2} S_{r_1} \mathbf{X},$$

where S_{r_i} is the graph shift operator of the edge-type subgraph $G_{r_i} = (V, A_{r_i}, X)$, for $i = 1, \dots, l$. Notably, given $S_{r_i} = A_{r_i}$, the multiplication $A_{r_l} \cdots A_{r_2} A_{r_1}$ is the diffusion process of a meta-path graph G_P based on P . The flexible choice of graph shift operator in the design of spectral GNNs allows us to view message propagation and aggregation within spatial domains as a special case of the spectral perspective, bridging these two domains and providing an intuitive grasp of multi-graph filters.

3.4 Multi-Graph Filters

A multi-graph spectral GNN [3, 17] is formulated as a multivariate polynomial $H(S_1, S_2, \dots, S_m)$, where S_1, \dots, S_m are graph shift operators, either adjacency matrices or Laplacian matrices of m homogeneous graphs G_1, \dots, G_m . Formally, H is defined as:

$$H(S_1, \dots, S_m) = \sum_{i=1}^m f_i(S_i) + f_{mv}(S_1, \dots, S_m), \quad (2)$$

where the first part processes signals within each graph G_i independently using univariate polynomial f_i , $i = 1, \dots, m$, while the second part f_{mv} incorporates multivariate monomials, e.g. $S_3 S_2 S_1$. This multi-graph filter degenerates to an ordinary spectral GNN for homogeneous graphs when $m = 1$.

3.5 High-frequency Area

Spectral GNNs leverage the Laplacian spectrum of input graphs as graph shift operators to collect spectral information: For any given graph $G = (V, A, X)$, the Laplacian spectrum can be represented by its Laplacian matrix $L = D - A$, where D is the diagonal degree matrix and A is the adjacency matrix. The spectral decomposition is defined as $L = U\Lambda U^\top$, where U contains normalized eigenvectors as columns and Λ is a diagonal matrix of eigenvalues. Following previous work [34], the high-frequency area can be computed by

$$S_{high} = \frac{\mathbf{x}^\top L \mathbf{x}}{\mathbf{x}^\top \mathbf{x}} \in \mathcal{R}^1, \quad (3)$$

where \mathbf{x} serves as a graph signal for nodes in G , and L is the graph Laplacian matrix. According to existing works [11, 14, 34], S_{high} reflects both the high-frequency content and the smoothness of graph signal \mathbf{x} from spectral and spatial domains, respectively, providing valuable properties to tackle challenges mentioned in Section 1. Moreover, its computation efficiency, as it avoids eigenvalue calculations [34], makes it an ideal metric for profiling the spectral energy distribution of signal \mathbf{x} on graph G .

3.6 Heterogenous Graph Anomaly Detection

Given a heterogeneous graph $G_M = (\mathcal{V}, \mathcal{A}, \mathcal{X}, O_V, R_E)$ and a target node type $o_t \in O_V$, let $V_{o_t}^{an}, V_{o_t}^{no} \subseteq V_{o_t}$ be two disjoint node subsets ($V_{o_t}^{an} \cap V_{o_t}^{no} = \emptyset$) representing anomalous and normal nodes, respectively. The GAD task aims to identify nodes belonging to $V_{o_t}^{an}$, which typically exhibit atypical patterns regarding structures or attributes [33]. In addition, unlike general binary classification, GAD faces a severe class imbalance challenge, where $|V_{o_t}^{an}| \ll |V_{o_t}^{no}|$

[11, 33]. Building on the foregoing discussion, this study aims to devise a novel spectral GNN to address the key challenges outlined in Section 1, specifically tailored for node-level GAD in heterogeneous networks. We present the detailed framework in Section 4.

4 ChiGAD

Wavelet filters have shown superiority in various tasks [29, 40] including GAD [2, 28, 34], which motivates our proposal of the Chi-Square filter. In this section, we first introduce the Chi-Square filter, which serves as the fundamental component of our ChiGAD framework, proving its wavelet properties and highlighting its unique characteristics. Based on these preliminaries, we present ChiGAD as a comprehensive solution to the three challenges outlined in Section 1. ChiGAD comprises three key components: (1) Multi-Graph Chi-Square Filter, (2) Interactive Meta-Graph Convolution, and (3) Contribution-Informed Cross-Entropy Loss. The overall architecture of ChiGAD is depicted in Figure 1.

4.1 Wavelet Properties of Chi-Square

Recent research has demonstrated the remarkable efficacy of graph wavelet [16] among various graph filters [3, 5, 17, 18, 34], particularly in domains such as anomaly detection, signal processing and computer vision [2, 28, 29, 34, 40]. Inspired by these studies, we introduce a novel graph wavelet based on the Chi-Square distribution. In this section, we demonstrate that the Chi-Square distribution not only satisfies the theoretical requirements of a graph wavelet but also offers a unique additivity property. This additivity, complementing the established advantages of conventional wavelets, proves particularly valuable for heterogeneous GAD.

Consider the Chi-Square distribution with n degrees of freedom. Its probability density function is defined as

$$f(x) = \frac{1}{2^{\frac{n}{2}} \Gamma(\frac{n}{2})} x^{\frac{n}{2}-1} e^{-\frac{x}{2}}, \quad x > 0,$$

where e is the natural logarithm base and $\Gamma(\cdot)$ is the Gamma function [8]. While $f(x)$ is not a linear polynomial, it can be well approximated by 3-order Chebyshev polynomials [19, 35], as validated by empirical results. In this work, we introduce a scaling factor m , and the Chi-Square filter with parameter i is defined as:

$$\begin{aligned} f_i(w) &= \frac{1}{S_i} \frac{1}{2^{\frac{n}{2}} \Gamma(\frac{n}{2})} \left(\frac{w}{m}\right)^{\frac{n}{2}-1} e^{-\frac{w}{2m}} \\ &= \frac{1}{S_i} \frac{1}{2^i \Gamma(i)} (w(i+1))^{i-1} e^{-\frac{w(i+1)}{2}}, \quad i = 1, 2, \dots, \end{aligned} \quad (4)$$

where $n = 2i, m = \frac{1}{i+1}$, and $S_i = \int_0^2 \frac{1}{2^i \Gamma(i)} (w(i+1))^{i-1} e^{-\frac{w(i+1)}{2}} dw$ regularizes f_i into a probability density function in $[0, 2]$, which is the full frequency range [34]. Such scaling enhances the expressiveness of f_i in processing signals of different frequencies while retaining the crucial properties of f_i , which are detailed below:

(1) Admissibility Condition:

$$\int_0^\infty \frac{|f_i(w)|^2}{w} dw < \infty.$$

It ensures filter f_i satisfies wavelet requirements and maintains band-pass property [16]. Proof is shown in Appendix C.1.

(2) Broad Band Coverage: Ideal filters $\{f_i\}_{i=1}^\infty$ are expected to cover diverse frequency bands. To this end, empirical results shown in Appendix C.2 verify the broad coverage of both $\{\mathbb{E}(\mathbf{X}_i) \mid \mathbb{E}(\mathbf{X}_i) = \frac{1}{S_i} \frac{m}{n} = \frac{1}{S_i} \frac{n}{m} = \frac{1}{S_i} \frac{2i}{i+1}\}_{i \geq 1}$ and $\{f_i^{-1}(f_i^{\max}) \mid f_i^{-1}(f_i^{\max}) = \arg \max_x f_i(x)\}_{i \geq 1}$, which are the expectation of $\mathbf{X}_i \sim f_i(x)$ and the most highlighted frequency for any $i \geq 1$, respectively.

(3) Spectral Locality: For $\mathbf{X}_i \sim f_i(x)$, $\forall i = 2, 3, \dots$,

$$\mathbb{E}(\mathbf{X}_i) = \frac{1}{S_i} \frac{2i}{i+1}, \quad \mathbb{V}(\mathbf{X}_i) = \frac{1}{S_i} \frac{2i}{(i+1)^2}.$$

As f_i is regularized to a distribution in $[0, 2]$, we have $\mathbb{E}(\mathbf{X}_i) \in [0, 2]$ and $\mathbb{V}(\mathbf{X}_i) \rightarrow 0$ when $i \rightarrow \infty$.

(4) Spatial Locality: When using a d -order Chebyshev polynomial, $f_i(L)$ is an $(i-1+d)$ -order polynomial of the Laplacian matrix L . Following Lemma 5.2 in existing research [16], signal effects are localized within $(i-1+d)$ -hop neighbors (where $d = 3$).

4.2 Multi-Graph Chi-Square Filter

Meta-Path Graph Generation. Node heterogeneity, characterized by diverse node types, presents several challenges for graph representation learning, such as attribute dimension disparity [43]. Our approach addresses this by learning representations for each node type independently using Chi-Square filters, a strategy that not only mitigates dimensional disparity but also enables effective representation capture, leveraging the powerful properties of the Chi-Square filter, as shown in our experiments.

To capture rich semantic relationships, we first generate a set of meta-paths P , as defined in Equation 1, ensuring that:

$$o_1 = o_{l+1} = o,$$

i.e., any meta-path is selected if its length is within a specified range and its first and last nodes are of the same type. This condition ensures homogeneity in the resulting meta-path graphs for each node type $o \in O_V$, denoted as $P_o^1, P_o^2, \dots, P_o^{k_o}$. The process is shown on the left side of Figure 1(a).

Empirically, the frequency focuses of these diverse meta-path graphs are distinct from each other, i.e., their spectral energy distributions exhibit different patterns. This observation motivates us to assign dedicated Chi-Square filter to each meta-path graph, ensuring that the frequency focus of the filter aligns with that of the graph. This approach has proven effective in homogeneous GAD scenarios [5, 14, 34]. Based on this, we detail this process in three steps. First, we derive the frequency focus of both meta-path graphs (step 1) and Chi-Square filters (step 2). Then, we present the filter assignment strategy to generate final representations (step 3).

Step 1: Meta-Path Graph Frequency Focus Estimation. Although the spectral energy [14, 34] has comprehensively characterized each meta-path graph from the spectral domain, it involves eigen-related computation, which is highly time-consuming. Thus, we do not calculate the energy distribution for all meta-path graphs. Instead, we use S_{high} as mentioned in Section 3.5, which has been shown to reflect the frequency characteristics of the graph effectively [14, 34], to select a small number of representative meta-path graphs for the aforementioned computation. Particularly, we divide all the valid meta-path graphs into three divisions of equal size by ranking on S_{high} , identifying them as low, mid, and high divisions,

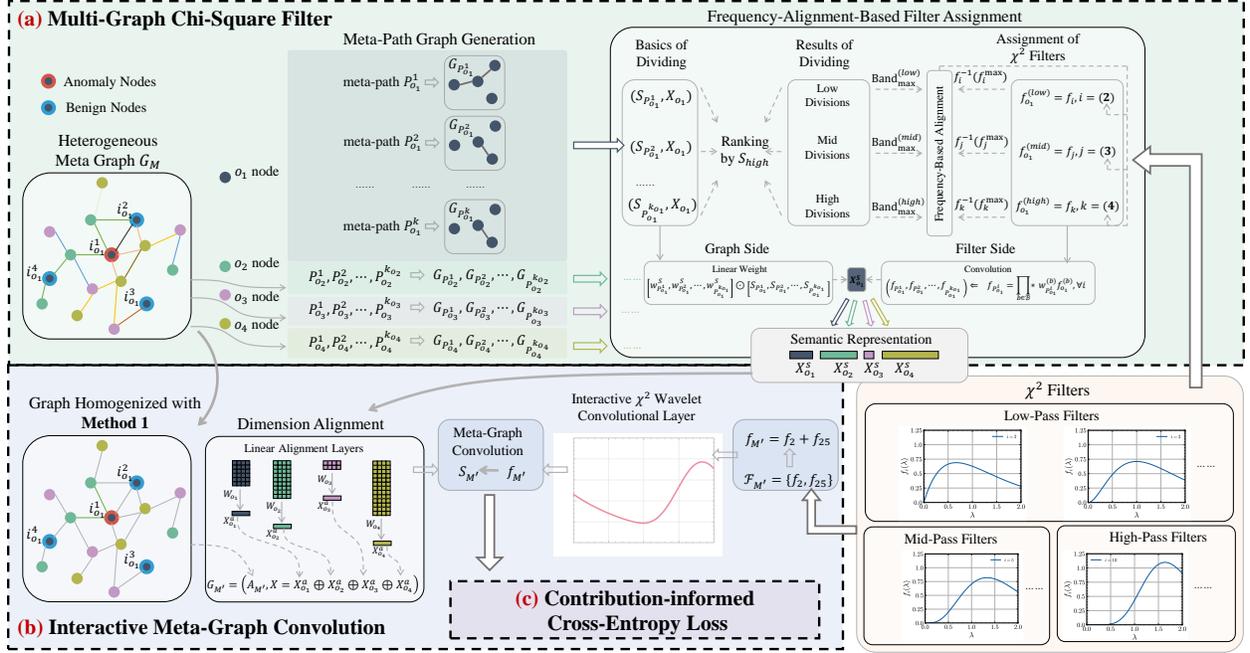


Figure 1: Illustration of the ChiGAD framework, comprising three modules: (a) Multi-Graph Chi-Square Filter; (b) Interactive Meta-Graph Convolution; (c) Contribution-Informed Cross-Entropy Loss.

respectively. Then, based on this division and the proven close connection between S_{high} and spectral energy [14, 34], the meta-path graph with median S_{high} in each division is regarded as the representative of this division, i.e. its spectral energy is representative of other meta-path graphs in the division. They can be denoted as $G_{P_o}^{(b)}$ for node type $o \in O_V$ and $b \in \{low, mid, high\}$. Thereby, we then fine-grainedly examine their frequency focuses. Denote the eigenvalues (sorted in ascending order), eigenvectors, Laplacian matrix of representative meta-path graph $G_{P_o}^{(b)}$ and the k -th column of attribute matrix as $\{\lambda_j^{(b)}\}_{j=1}^{|V_o|}, \mathbf{U}_{[:,j]}^{(b)}, L_o^{(b)}$ and \mathbf{x}_o^k , respectively, where $k = 1, \dots, d_o$. The Fourier coefficients are expressed as:

$$\tilde{\mathbf{x}}_{o,j}^{(b)} = \sum_{k=1}^{|d_o|} \left(\mathbf{U}_{[:,j]}^{(b)} \right)^T \mathbf{x}_o^k, \quad j = 1, 2, \dots, |V_o|, \quad (5)$$

which capture the spectral energy of the graph $G_{P_o}^{(b)}$. The energy associated with each eigenvalue $\lambda_j^{(b)}$ is defined as $w_j^{(b)} = \left(\tilde{\mathbf{x}}_{o,j}^{(b)} \right)^2$. Rather than considering the contribution of each eigenvalue individually, we exploit the excellent band-pass property of the Chi-Square filter by analyzing the focus on a band-by-band basis. To this end, we partition the full frequency band $[0, 2]$ into K sequential, disjoint bands. Formally, for each band k , we define:

$$\text{Band}_k = \text{List} \left[\lambda_j^{(b)} \right]_{j=k}^{k + \frac{|V_o|}{K} - 1}, \quad w_{\text{Band}_k}^{(b)} = \sum_{j=k}^{k + \frac{|V_o|}{K} - 1} w_j^{(b)}, \quad (6)$$

with $k \in \{1, 1 + \frac{|V_o|}{K}, 1 + 2 \frac{|V_o|}{K}, \dots, |V_o| - \frac{|V_o|}{K} + 1\}$. The band with the maximum cumulative weight is then identified as

$$\text{Band}_{\max}^{(b)} = \text{REP} \left(\arg \max_{\text{Band}_k} w_{\text{Band}_k}^{(b)} \right) \in \mathcal{R}^1, \quad (7)$$

where REP denotes a representative operation on Band_k (e.g., computing the median or mean). Consequently, the spectral focus for all meta-path graphs in division b is defined to be $\text{Band}_{\max}^{(b)}$.

Step 2: Obtaining Frequency Focus for Chi-Square Filters. Unlike the complex band analysis used for meta-path graphs, the frequency focus of each Chi-Square filter f_i ($i = 1, 2, \dots$) can be defined more intuitively. Specifically, we identify the most prominent frequency as $f_i^{-1}(f_i^{\max}) \in \mathbb{R}^1$, where f_i^{\max} is the maximum value of $f_i(w)$ (defined in Equation 4), and f_i^{-1} is its inverse function.

Step 3: Assigning Filter and Semantic Representation Learning. Since now both the frequency focuses of meta-path graphs and Chi-Square filters are available, it is now feasible to obtain the Chi-Square filter $f_o^{(b)}$ that mostly aligns with the frequency focus $\text{Band}_{\max}^{(b)}$ of $G_{P_o}^{(b)}$ by

$$f_o^{(b)} = \arg \min_{f_i} |f_i^{-1}(f_i^{\max}) - \text{Band}_{\max}^{(b)}| \quad (8)$$

and subsequently learn the node representation across diverse meta-path graphs with assigned Chi-Square filters. Prior to the final representation learning, the further process for meta-path graphs and assigned filters are conducted, respectively. For meta-path graph side, weight $w_{P_o}^S$ is assigned to the graph shift operator $\mathcal{S}_{P_o}^i$ of each meta-path graph $G_{P_o}^i$ to learn the distinct importance of diverse meta-paths; for Chi-Square filter side, based on the additivity of the Chi-Square distribution, we propose a novel fusion method Chi-Square filter, which is formally defined as:

$$f_{P_o}^i = \left(\prod_{b \in \mathcal{B}} * \right) w_{P_o}^{(b)} f_o^{(b)},$$

where $f_{p_o^i}$ is the filter assigned for each specific meta-path graph $G_{p_o^i}$, \prod * is continuous convolution between filters, $\mathcal{B} = \{low, mid, high\}$, and $w_{p_o^i}^{(b)}$ is a division-wise weight defined as

$$w_{p_o^i}^{(b)} = \begin{cases} 1, & \text{if } G_{p_o^i} \text{ is ranked as a } (b)\text{-band-focused graph} \\ w_d, & \text{if } G_{p_o^i} \text{ is not ranked as a } (b)\text{-band-focused graph} \end{cases}$$

to retain the focus of the convolved filter $f_{p_o^i}$ in the bands highlighted by the spectral energy of $G_{p_o^i}$ while introducing neighboring bands as supplements. $w_d \ll 1$ serves as a hyper-parameter. To make such a form of convolution clearer.

In summary, following the definition of multi-graph filter in Equation 2, the Multi-Graph Chi-Square Filter H_o formally processes the input representation X_o as:

$$\begin{aligned} X_o^s &= \text{Sum} \left(\left[w_{p_o^1}^s, \dots, w_{p_o^{k_o}}^s \right] \odot \left[S_{p_o^1}, \dots, S_{p_o^{k_o}} \right] \odot \left[f_{p_o^1}, \dots, f_{p_o^{k_o}} \right] \right) \\ &= \sum_{i=1}^{k_o} f_{p_o^i} \left(w_{p_o^i}^s S_{p_o^i} \right) X_o = H_o(S_{p_o^1}, \dots, S_{p_o^{k_o}}) X_o \end{aligned} \quad (9)$$

for any $o \in O_V$, where σ is an activation function. X_o^s serves as **semantic representation** of nodes of type o , where semantics across various meta-paths and diverse frequency bands are integrated. In this process, the additivity property of the Chi-Square plays an essential role in two aspects:

- The additivity of Chi-Square distribution ensures that each monomial $f_{p_o^i}$ within the convolved filter H_o still aligns with the definition of graph wavelet, retaining excellent properties detailed in Section 4.1. Thus, the representation X_o^s is a linear combination of the learned outputs from multiple Chi-Square filters.
- The additivity of Chi-Square promotes collaborative representation learning across diverse meta-paths in an effective way. Specifically, it not only ensures the graph with signals are processed by the meticulously assigned and most aligned filter as mentioned but also incorporates additional filters naturally derived from convolutional fusion to enhance broader frequency-band coverage.

4.3 Interactive Meta-Graph Convolution

Processing each node type simultaneously leaves dimensional disparities unresolved—a major challenge in heterogeneous settings. Moreover, retaining high-frequency information is crucial for GAD tasks, as it helps counteract over-smoothing [5, 14, 34]. Yet, the ability of alignment methods to preserve high-frequency content during dimensional alignment in HINs remains largely unexplored.

In this study, we propose to use the linear layer to align feature dimensions. Despite its simplicity, our provided theoretical analysis in this section and empirical results in Appendix 5.4 have jointly proven that the linear layer is a nice choice as the alignment component of the GAD model in heterogeneous scenario by effectively preserving high-frequency contents. Next, we first introduce an intuitive revelation of the linear layers' functionality.

An Intuitive Revelation of the Functionality. For nodes V_o of any type $o \in O_V$, the semantic representation matrix learned from Equation 9 is comprised of d_o^s signals as column vectors, which is

formally expressed as:

$$X_o^s = x_o^{s1} \oplus x_o^{s2} \oplus \dots \oplus x_o^{sd_o^s} \in \mathcal{R}^{|V_o| \times d_o^s}.$$

Given a node-type-wise linear layer $W_o = (w_{ij})_{d_o^s \times d_a} \in \mathcal{R}^{d_o^s \times d_a}$ where d_a represents the aligned dimension, the aligned representation is computed by:

$$X_o^a = X_o^s W_o = x_o^{a1} \oplus x_o^{a2} \oplus \dots \oplus x_o^{ad_a} \in \mathcal{R}^{|V_o| \times d_a}. \quad (10)$$

Equation 10 reveals the derivation of each aligned representation $x_o^{a,j}$ is a linear summation of the raw graph signals as:

$$x_o^{a,j} = \sum_{i=1}^{d_o^s} w_{ij} x_o^{si}, \quad j = 1, 2, \dots, d_a,$$

which indicates the linear layer W_o learns each new graph signal $x_o^{a,j}$, $j = 1, \dots, d_a$ by automatically adjusting the relationships between the raw graph signals $\{x_o^{si}\}_{i=1}^{d_o^s}$ with learnable linear weights.

Theoretical Justification of Linear Summation. The vital importance of counteraction against over-smoothness and retaining high-frequency to GAD task have been deeply investigated in the latest researches [5, 14, 34]; meanwhile, S_{high} serves as a handy depiction for the smoothness and high-frequency content of any graph signal from the spatial and spectral domains, respectively [14, 34]. Based on these two findings, we then prove the linear layer's excellent capability of high-frequency retention.

Theorem 1. *For a given collection of k graph signals x_1, \dots, x_k defined on a graph with Laplacian matrix L . In theory, there exists a linear alignment layer such that the aligned node representation x_a perfectly keeps the S_{high} . Formally, $S_{high}(x_a, L) = \max \{S_{high}(x_i, L)\}_{i=1}^k$.*

Proof of Theorem 1 is presented in Appendix C.3. Theorem 1 illuminates the accessibility of optimal alignment with linear layers when the preservation of S_{high} is taken as criterion. Notably, the fact that the linear transformations represented by W_o are not constant is not a detriment to performance in both retention and GAD tasks; on the contrary, this flexibility is shown to be beneficial in experimental results of Appendix 5.4, probably because it enables W_o to automatically learn the linear mapping that is most effective for anomaly detection; in addition, other alignment methods with similar linear shapes, e.g. PCA, AutoEncoder also show significant performance improvements over nonlinear methods in both high-frequency retention and GAD tasks in Appendix 5.4, which further verifies our theoretical findings. These subsequent theoretical and empirical results motivate us to propose to align the representations of diverse nodes as:

$$X_o^a = X_o^s W_o, \quad o = 1, 2, \dots, |O_V|, \quad (11)$$

where $W_o \in \mathcal{R}^{d_o^s \times d_a}$ is a node-type-wise linear layer and X_o^a is the aligned representation of type o nodes. Then, employing Method 1 as detailed in Section 5.1, we derive the transformed homogeneous graph G'_M comprising all types of nodes, and thence learn the final node representations formally as:

$$X'_o = \sum_{f \in \mathcal{F}_{M'}} f(S_{M'}) \sigma(X_o^a), \quad o = 1, 2, \dots, |O_V|, \quad (12)$$

where $\mathcal{F}_{M'}$ is a specified set of Chi-Square filters covering diverse frequency bands and $S_{M'}$ is the graph shift operator of $G_{M'}$.

4.4 Contribution-Informed Cross-Entropy Loss

Prior works [11, 14] have highlighted the significant challenges posed by the class imbalance inherent in anomaly graphs within GAD. This motivates us to develop this module to enhance the effective learning of anomalies. Specifically, in this module, we firstly obtain a homogeneous graph G_{o_t} for the target node type o_t using Method 2 delineated in Section 5.1. For the final representation $\mathbf{X}'_{o_t} = \mathbf{x}'_{o_t,1} \oplus \mathbf{x}'_{o_t,2} \oplus \dots \oplus \mathbf{x}'_{o_t,d'_{o_t}} \in R^{|V_{o_t}| \times d'_{o_t}}$ and the Laplacian matrix \mathbf{L}_{o_t} of G_{o_t} , the contribution of each node is defined as:

$$c_i = \sum_{j=1}^d \frac{\mathbf{x}'_{o_t,j,i} (\mathbf{L}_{o_t} \mathbf{x}'_{o_t,j,i})^{j,i}}{\mathbf{x}'_{o_t,j,i}{}^T \mathbf{L}_{o_t} \mathbf{x}'_{o_t,j,i}}, \quad i = 1, \dots, |V_{o_t}|,$$

where $\mathbf{x}'_{o_t,j,i}, (\mathbf{L}_{o_t} \mathbf{x}'_{o_t,j,i})^{j,i} \in \mathcal{R}^1$ are the i -th element of column vectors $\mathbf{x}'_{o_t,j}$ and $\mathbf{L}_{o_t} \mathbf{x}'_{o_t,j}$, respectively. With a focus on the distribution of contributions, the following empirical observations can be revealed:

- In most cases, the average contribution of anomalies is significantly larger than that of benign nodes within the train set;
- In most cases, the average contribution of recognized anomalies (i.e. true positive/TP) is significantly larger than that of unrecognized anomalies (i.e. false negative/FN) within train set, which jointly indicate that the benign and FN anomalies within the train set exhibit similar overall deficiencies in terms of contributions to S_{high} . Such similarity is not surprising, since the representations of FN anomalies are generally more challenging to distinguish from benign nodes than those of TP anomalies.

Building upon the crucial findings and analyses, we propose the Contribution-Informed Cross-Entropy Loss, formally defined as:

$$\mathcal{L}_{CC} = -\frac{1}{N} \sum_{i=1}^N w_i \left[y_i \log(\hat{p}_i) + (1 - y_i) \log(1 - \hat{p}_i) \right],$$

where for any i -th sample in the train set, y_i is the true label,

$$\hat{p}_i = \text{softmax}(\mathbf{X}'_{o_t})_i$$

is the predicted anomaly probability, and

$$w_i = \begin{cases} 1, & \text{if } i\text{-th sample} \in V^{bn}, \\ \frac{c_{\max} - c_i}{c_{\max} - c_{\min}} (H - L) + L, & \text{if } i\text{-th sample} \in V^{an}, \end{cases}$$

is the contribution-informed weight. H, L are hyper-parameters to regulate the range of weights, and c_{\min}, c_{\max} are computed by:

$$c_{\min} = \min\{c_i\}_{i=1}^N, \quad c_{\max} = \max\{c_i\}_{i=1}^N.$$

Notably, we impose an essential constraint that $H \geq L \geq 1$, which ensures that the model prioritizes challenging anomalies (FN) and gives greater attention to anomalous samples, which are of scarcity, compared to benign nodes. As the weights assigned to anomalous samples are ensured to be greater, this constraint helps mitigate the inherent class-imbalance problem in anomaly detection tasks.

5 Experiments

In this section, we conduct holistic experiments to evaluate the capability of ChiGAD in heterogeneous graph anomaly detection.

5.1 Experimental Setup

Datasets. We assess ChiGAD using one public and two industrial datasets, and further evaluate ChiGNN, a ChiGAD-derived homogeneous graph anomaly detector, on seven public datasets:

- **Heterogeneous graph for ChiGAD:** Public ACM dataset and two real-world datasets R-I and R-II provided by WeChat Pay, Tencent, which pertain to financial activities in two different regions.
- **Homogeneous graph for ChiGNN:** Amazon, T-Finance, T-Social, Reddit, Tolokers, Questions, and DGraph-Fin.

The summaries of datasets are detailed in Appendix A. It is worth noting that ACM is not a natural dataset for GAD tasks. However, the node of type o_0 has 3 classes, which allows us to fix the target node type $o_t = o_0$ and randomly select one of the classes as the anomalous class and the other two as the benign class.

Baselines and Implementation Details. We introduce a diverse range of cutting-edge baselines for different types of experiments to compare and evaluate the model performance:

- **Baselines for ChiGAD:** Operational-Model, PSHGCN [17], BWGNN [34], XGBGraph [33], and ConsisGAD [6];
- **Baselines for ChiGNN:** GAT [36], GIN [41], PCGNN [26], BernNet [18], PSHGCN [17], AMNet [5], BWGNN [34], GAGA [39], GHRN [14], XGBGraph [33], and ConsisGAD [6].

Note that, for fair comparison in all experiments, we obtain the source code of these baselines from GitHub and use the default hyper-parameter setting recommended by their authors. In addition, to ensure compatibility with models not originally designed for heterogeneous graphs, we degenerate heterogeneous graphs, ACM, R-I, and R-II, into homogeneous graphs through two distinct methods, and adopt the **superior** performance obtained from these two methods for a fair comparison:

- **Method 1:** All the nodes within the meta graph, i.e. the raw heterogeneous graph, are reserved; any edge is reserved if and only if it exists in at least one relation. In addition, the problem of dimensional disparity is addressed initially with linear layer;
- **Method 2:** We keep only the nodes of target type o_t . The inter-connection is formed only when the endpoints are connected in the context of any considered relation.

As for the implementation setup, we employ the fully-supervised settings following the up-to-date benchmark GADBench [33]. Detailed information for reproducibility is in Appendix B.1 and B.2.

Metrics. In this study, **AUROC**, **AUPRC**, **F1-macro**, and **Recall** are selected for providing reliable and holistic evaluation, since the former two are threshold-free while the rest are threshold-based.

5.2 Performance Comparison

Table 1 shows the performance of ChiGAD against diverse baselines across 3 heterogeneous graphs. We highlight the best performance among competitors in bold text.

Firstly, ChiGAD significantly surpasses all the state-of-the-art methods in all heterogeneous datasets across all metrics, demonstrating its superiority in heterogeneous graph anomaly detection.

Secondly, the explicit capture of rich semantics across diverse meta-path graphs has been demonstrated to be crucial for heterogeneous graph anomaly detection. On one hand, ChiGAD and

Table 1: Experimental results of performance comparison on heterogeneous datasets (ACM, R-I, R-II).

Datasets Metrics	ACM				R-I				R-II			
	AUROC	AUPRC	F1-macro	Recall	AUROC	AUPRC	F1-macro	Recall	AUROC	AUPRC	F1-macro	Recall
Operational-Model	-	-	-	-	0.8200	0.0473	0.5528	0.1105	0.8120	0.0482	0.5520	0.1378
PSHGCN	0.9564	0.6226	0.7992	0.7305	0.8650	0.2264	0.6425	0.3316	0.9122	0.4355	0.7148	0.4551
BWGNN	0.9030	0.8452	0.4015	0.7444	0.8444	0.2156	0.6491	0.3263	0.9133	0.2231	0.6197	0.2231
XGBGraph	0.9461	0.9068	0.8647	0.8271	0.8795	0.2960	0.6812	0.3632	0.9155	0.5696	0.8014	0.5992
ConsisGAD	0.8773	0.7462	0.8066	0.7461	0.8724	0.2944	0.6716	0.3747	0.9014	0.4500	0.7284	0.4696
ChiGAD	0.9702	0.9518	0.9195	0.8947	0.9472	0.3887	0.7272	0.4579	0.9322	0.5854	0.8103	0.6133
Improvement	2.55%	4.96%	6.34%	8.17%	7.70%	31.34%	6.75%	26.09%	1.82%	2.77%	1.11%	2.37%

Table 2: Experimental results of performance comparison on homogeneous datasets.

Datasets	Metrics	GAT	GIN	PCGNN	BernNet	PSHGCN	AMNet	BWGNN	GAGA	GHRN	XGBGraph	ConsisGAD	ChiGNN
Reddit	AUROC	0.6415	0.6182	0.6541	0.6983	0.6744	0.6960	0.7082	0.5352	0.6102	0.6474	0.5859	0.7288
	AUPRC	0.0720	0.0641	0.0773	0.0782	0.0625	0.0787	0.0832	0.0359	0.0466	0.0529	0.0607	0.0953
	F1-macro	0.4915	0.4915	0.4915	0.5406	0.4915	0.4915	0.5129	0.4915	0.5149	0.5286	0.4915	0.5596
	Recall	0.1088	0.1088	0.1497	0.1088	0.0952	0.1156	0.1156	0.0204	0.0544	0.0612	0	0.1565
Tolokers	AUROC	0.7875	0.7457	0.7663	0.7751	0.7784	0.7516	0.8041	0.5001	0.8008	0.8285	0.7972	0.8349
	AUPRC	0.4525	0.4036	0.4485	0.4369	0.4494	0.4074	0.4958	0.2184	0.4745	0.5392	0.4691	0.5413
	F1-macro	0.6356	0.6087	0.6047	0.6152	0.5568	0.6200	0.6574	0.4387	0.6484	0.6688	0.6741	0.7137
	Recall	0.4424	0.3941	0.4315	0.4470	0.4517	0.4252	0.5031	0.2150	0.4657	0.5343	0.5252	0.5514
Amazon	AUROC	0.9713	0.9563	0.9801	0.9579	0.9726	0.9692	0.9827	0.9731	0.9829	0.9874	0.9823	0.9878
	AUPRC	0.8794	0.8461	0.8933	0.8489	0.8855	0.8836	0.9148	0.8807	0.8952	0.9333	0.9273	0.9160
	F1-macro	0.9034	0.2608	0.8997	0.8941	0.9072	0.8956	0.9034	0.9155	0.9039	0.9276	0.9265	0.9300
	Recall	0.8261	0.8098	0.8533	0.8261	0.8533	0.8315	0.8587	0.8376	0.8533	0.8533	0.8061	0.8859
T-Finance	AUROC	0.9575	0.9271	0.9403	0.9667	0.9735	0.9638	0.9693	0.9439	0.9646	0.9715	0.9736	0.9749
	AUPRC	0.8272	0.7835	0.8331	0.8917	0.8961	0.8887	0.8938	0.8079	0.8760	0.9012	0.8937	0.9029
	F1-macro	0.7296	0.8254	0.7949	0.7221	0.9141	0.8273	0.6505	0.8041	0.6484	0.9308	0.9191	0.9319
	Recall	0.7975	0.7337	0.7906	0.8363	0.8377	0.8405	0.8419	0.7916	0.8197	0.8502	0.8278	0.8544
Questions	AUROC	0.6830	0.6808	0.6766	0.6858	0.6935	0.6605	0.7087	0.5003	0.7216	0.7102	0.7316	0.7373
	AUPRC	0.1551	0.1368	0.1559	0.1725	0.1394	0.1563	0.1857	0.0299	0.1831	0.1819	0.1797	0.1930
	F1-macro	0.5726	0.5709	0.5799	0.5778	0.5572	0.5780	0.5878	0.4924	0.5857	0.5879	0.5896	0.6057
	Recall	0.1726	0.1836	0.1808	0.1973	0.1753	0.1753	0.2164	0.0301	0.2219	0.2055	0.1340	0.2329
DGraph-Fin	AUROC	0.7553	0.7416	0.7276	0.7358	0.7010	0.7299	0.7630	0.5000	0.7613	0.7583	0.7040	0.7640
	AUPRC	0.0385	0.0347	0.0342	0.0327	0.0335	0.0281	0.0397	0.0127	0.0380	0.0379	0.0080	0.0403
	F1-macro	0.5317	0.5150	0.5072	0.5199	0.4968	0.5177	0.5309	0.4968	0.5323	0.5259	0.4989	0.5344
	Recall	0.0714	0.0632	0.0666	0.0555	0.0469	0.0421	0.0757	0.0125	0.0696	0.0696	0	0.0800
T-Social	AUROC	0.9033	0.9409	0.8769	0.9373	0.9184	0.9250	0.9688	0.9696	0.9712	0.9976	0.9636	0.9980
	AUPRC	0.3207	0.6079	0.8029	0.4430	0.7111	0.3770	0.7893	0.8345	0.8678	0.9734	0.7213	0.9833
	F1-macro	0.5740	0.5462	0.4801	0.4341	0.8241	0.4879	0.6304	0.7319	0.6455	0.9256	0.8224	0.9760
	Recall	0.4207	0.6447	0.7353	0.4823	0.7034	0.4321	0.7578	0.7393	0.8233	0.9353	0.6089	0.9535

PSHGNCN, as multi-graph filters, both explicitly introduce various graph shift operators to capture semantics based on various meta-paths, and both show excellent GAD capabilities on the three datasets (it is worth noting that PSHGNCN is not designed for GAD tasks); on the other hand, in contrast, despite being state-of-the-art methods in the field of general homogeneous GAD, ConsisGAD and BWGNN perform worse than PSHGNCN due to their lack of multi-graph processing capabilities.

Then, we further compare ChiGAD with the two models designed for heterogeneous graph learning, i.e. Operational-Model and PSHGNCN. (1) Operational-Model is a model actually applied in the business of company R. It is based on GCN and learns features on multiple manually selected important meta-path graphs. Due to privacy reasons, Operational-Model cannot be obtained for offline deployment on the ACM dataset. Empirically, it shows the worst performance on the business datasets R-I and R-II. This can be attributed to the low-pass property of GCN, which causes a large amount of information loss when facing input signals containing diverse frequency band information, which is crucial for GAD [34].

(2) Although PSHGNCN is also a multi-graph filter, it is significantly inferior to ChiGAD across all the datasets and all the metrics. This proves that our proposed key components of ChiGAD designed for GAD help detect anomalies in heterogeneous graphs effectively.

ChiGAD significantly outperforms ConsisGAD, BWGNN, and XGBGraph in heterogeneous graph anomaly detection, as evidenced by superior performance across AUROC, AUPRC, F1-macro, and Recall on ACM, R-I, and R-II datasets. While ConsisGAD addresses limited supervision through consistency training, it lacks explicit handling of heterogeneity. BWGNN, designed for homogeneous graphs, struggles with the complexity of heterogeneous networks, and XGBGraph, despite its neighbor aggregation strength, fails to model fine-grained relational properties in heterogeneous graphs. In contrast, ChiGAD leverages multi-graph Chi-Square filters to capture cross-meta-path anomalies, preserves high-frequency content via linear alignment, and mitigates class imbalance with contribution-informed loss, showcasing the inherent advantages of multi-graph filters in heterogeneous anomaly detection.

Table 3: Comparison of results of GAD task across diverse alignment methods on ACM.

Metric	AUROC	AUPRC	F1-macro	Recall
LL (employed)	0.9702	0.9518	0.9195	0.8947
<i>PCA</i>	0.6587	0.5035	0.5975	0.4687
<i>UMAP</i>	0.5746	0.3852	0.5393	0.3885
<i>IsoMAP</i>	0.5714	0.3854	0.4417	0.3784
<i>AE</i>	0.9597	0.9088	0.9157	0.8847
<i>AE_10</i>	0.9483	0.8928	0.8798	0.8296
<i>AE_100</i>	0.9317	0.8782	0.8511	0.7945
<i>AE_1000</i>	0.9072	0.8517	0.8193	0.7845
<i>AE_10000</i>	0.8450	0.7510	0.7536	0.6667

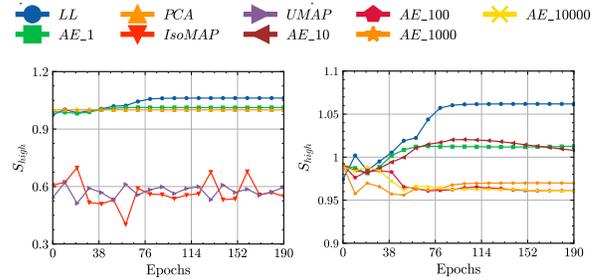
5.3 Superiority of Chi-Square Filter

In this study, the Chi-Square filter serves as the foundation for the multi-graph filter. Furthermore, as shown in Equation 12, in scenarios where dimensional disparities are not a concern, the Chi-Square convolutional layers, namely **ChiGNN**, can be separately compatible for homogeneous graph anomaly detection tasks. To investigate this, we examined the performance of as many as seven public datasets, with the experimental results presented in Table 2. It is evident that the ChiGNN outperforms all baselines in GAD tasks. This can be attributed to the excellent properties with the Chi-Square filter, e.g., both the spatial and spectral locality, as well as the wide coverage on diverse frequency bands. For reproducibility, the hyper-parameter settings are listed in Appendix B.2.

5.4 Resisting Over-Smoothness for GAD: Comparison of Alignment Methods

This section aims at comparing the performance of diverse alignment methods on both the frequency retention task and the GAD task, which further helps reveal the relationship between resisting over-smoothness and anomaly detection. To this end, we first compare the performance of linear layer (LL), which is employed as dimension alignment component of ChiGAD as detailed in Section 4.3 with diverse representative alignment methods, i.e. Autoencoder (AE), Principle Component Analysis (PCA), Uniform Manifold Approximation and Projection (UMAP), and Isometric Mapping (IsoMAP) on the frequency retention task. Based on this, we further compare the performance of ChiGAD and four variants with AE, PCA, UMAP, and IsoMAP as the alignment components on GAD tasks. It is worth noting that we take S_{high} as the metric to quantify both the smoothness content and high-frequency content retained (the frequency energy distribution of graph signal in the spectral domain, which is evaluated by S_{high} , is directly and closely related to its smoothness as node representation in the spatial domain [14, 34]. Specifically, higher S_{high} indicates a lower degree of smoothness and more high-frequency content is preserved from the spatial and spectral domains, respectively, and vice versa. [14, 34]); for the GAD task, AUROC, AUPRC, F1-macro and Recall are introduced. Experimental results are shown in Figure 2 and Table 3.

Figure 2 shows the S_{high} of aligned nodal representation X_o^a , which is defined in Equation 11, obtained in the training process. Notably, considering both the Laplacian matrices of any mentioned graph and raw nodal features X_o are constant (since the former

**Figure 2: Comparison of results of frequency retention task across diverse alignment methods on ACM.**

corresponds to structures of graphs, which are consistent), the initial high-frequency content evaluated by S_{high} is constant. Thus, a higher S_{high} of aligned representation indicates a better performance of any specific alignment method in frequency retention task. Based on this finding, subsequent observations are obtained:

- Linear layer achieves the optimal performance in frequency retention; furthermore, the performance of linear-related methods (LL, PCA, AE) significantly surpasses that of nonlinear methods (UMAP, IsoMAP). This empirical results highly align with our theoretical analysis in Section 4.3.
- Although AE is not fully linear due to activation functions, it resembles a stack of linear layers, explaining its comparable performance to LL and PCA.
- Among linear-related methods, LL is significantly superior to PCA and AE due to its flexibility in prioritizing high-frequency retention. Conversely, the inherent functionalities of PCA (variance maximization) and AE (reconstruction error minimization), despite also being linear-related, may conflict with this task.
- To further explore the above point, multiple variants of AE are introduced, where the reconstruction errors are amplified with weights from 10 to 10,000. As we can see, the variants with larger weights perform worse in the retention task, which indicates the task of reconstruction contradicts high-frequency retaining.

Results of the GAD task in Table 3 show similar patterns to those of the frequency retention task, which reveals the crucial connection between GAD and high-frequency retention. Linear methods outperform UMAP and IsoMAP, with LL being the best. AE outperforms PCA, especially with smaller reconstruction weights. This indicates that both PCA’s variance maximization and AE’s reconstruction conflict with GAD, causing greater information loss with PCA in aligning representation dimensions.

6 Conclusion

In this paper, we introduce ChiGAD, an effective framework for graph anomaly detection on heterogeneous networks using our proposed Chi-Square filter. Extensive experiments show the great efficacy of ChiGAD in heterogeneous GAD. Moreover, our ChiGNN variant surpasses state-of-the-art baselines on homogeneous GAD tasks, further validating the effectiveness of the Chi-Square filter.

Acknowledgments

This work is supported by the RGC GRF grant (No. 14217322) and the Tencent WeChat Rhino-Bird Focused Research Program.

References

- [1] Source Code. <https://github.com/HsipingLi/ChiGAD>.
- [2] Vicente Alarcon-Aquino and Javier A Barria. 2001. Anomaly detection in communication networks using wavelets. *IEE Proceedings-Communications* 148, 6 (2001), 355–362.
- [3] Landon Butler, Alejandro Parada-Mayorga, and Alejandro Ribeiro. 2023. Convolutional Learning on Multigraphs. *IEEE Trans. Signal Process.* 71 (2023), 933–946.
- [4] Iván Cantador, Peter Brusilovsky, and Tsvi Kuflik. 2011. Second workshop on information heterogeneity and fusion in recommender systems (HeRec2011). In *RecSys*. 387–388.
- [5] Ziwei Chai, Siqi You, Yang Yang, Shiliang Pu, Jiarong Xu, Haoyang Cai, and Weihao Jiang. 2022. Can abnormality be detected by graph neural networks?. In *IJCAI*. 1945–1951.
- [6] Nan Chen, Zemin Liu, Bryan Hooi, Bingsheng He, Rizal Fathony, Jun Hu, and Jia Chen. 2024. Consistency Training with Learnable Data Augmentation for Graph Anomaly Detection with Limited Supervision. In *ICLR*.
- [7] Darcy A. Davis, Ryan Lichtenwalter, and Nitesh V. Chawla. 2011. Multi-relational Link Prediction in Heterogeneous Information Networks. In *ASONAM*. 281–288.
- [8] Philip J Davis. 1959. Leonhard euler’s integral: A historical profile of the gamma function: In memoriam: Milton abramowitz. *The American Mathematical Monthly* 66, 10 (1959), 849–869.
- [9] Xiangyu Dong, Xingyi Zhang, Lei Chen, Mingxuan Yuan, and Sibow Wang. 2025. SpaceGNN: Multi-Space Graph Neural Network for Node Anomaly Detection with Extremely Limited Labels. In *ICLR*.
- [10] Xiangyu Dong, Xingyi Zhang, Yanni Sun, Lei Chen, Mingxuan Yuan, and Sibow Wang. 2025. SmoothGNN: Smoothing-aware GNN for Unsupervised Node Anomaly Detection. In *WWW*. 1225–1236.
- [11] Xiangyu Dong, Xingyi Zhang, and Sibow Wang. 2024. Rayleigh Quotient Graph Neural Networks for Graph-level Anomaly Detection. In *ICLR*.
- [12] Yuxiao Dong, Jie Tang, Sen Wu, Jilei Tian, Nitesh V. Chawla, Jinghai Rao, and Huanhuan Cao. 2012. Link Prediction and Recommendation across Heterogeneous Social Networks. In *ICDM*. 181–190.
- [13] Xinyu Fu, Jiani Zhang, Ziqiao Meng, and Irwin King. 2020. MAGNN: Metapath Aggregated Graph Neural Network for Heterogeneous Graph Embedding. In *WWW*. 2331–2341.
- [14] Yuan Gao, Xiang Wang, Xiangnan He, Zhengguang Liu, Huamin Feng, and Yongdong Zhang. 2023. Addressing Heterophily in Graph Anomaly Detection: A Perspective of Graph Spectrum. In *WWW*. 1528–1538.
- [15] William L. Hamilton, Zhitaoying, and Jure Leskovec. 2017. Inductive Representation Learning on Large Graphs. In *NeurIPS*. 1024–1034.
- [16] David K Hammond, Pierre Vandergheynst, and Rémi Gribonval. 2011. Wavelets on graphs via spectral graph theory. *Applied and Computational Harmonic Analysis* 30, 2 (2011), 129–150.
- [17] Mingguo He, Zhewei Wei, Shikun Feng, Zhengjie Huang, Weibin Li, Yu Sun, and Dianhai Yu. 2024. Spectral Heterogeneous Graph Convolutions via Positive Noncommutative Polynomials. In *WWW*. 685–696.
- [18] Mingguo He, Zhewei Wei, Zengfeng Huang, and Hongteng Xu. 2021. BernNet: Learning Arbitrary Graph Spectral Filters via Bernstein Approximation. In *NeurIPS*. 14239–14251.
- [19] Francis Begnaud Hildebrand. 1987. *Introduction to numerical analysis*. Courier Corporation.
- [20] Ziniu Hu, Yuxiao Dong, Kuansan Wang, and Yizhou Sun. 2020. Heterogeneous Graph Transformer. In *WWW*. 2704–2710.
- [21] Thomas N. Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *ICLR*.
- [22] Xiangnan Kong, Jiawei Zhang, and Philip S. Yu. 2013. Inferring anchor links across multiple heterogeneous social networks. In *CIKM*. 179–188.
- [23] Ji-chao Li, Dan-ling Zhao, Bing-Feng Ge, Ke-Wei Yang, and Ying-Wu Chen. 2018. A link prediction method for heterogeneous networks based on BP neural network. *Physica A: Statistical Mechanics and its Applications* 495 (2018), 1–17.
- [24] Xiping Li, Jianghong Ma, Kangzhe Liu, Shanshan Feng, Haijun Zhang, and Yutong Wang. 2024. Category-based and Popularity-guided Video Game Recommendation: A Balance-oriented Framework. In *WWW*. 3734–3744.
- [25] Fanzhen Liu, Zhao Li, Baokun Wang, Jia Wu, Jian Yang, Jiaming Huang, Yiqing Zhang, Weiqiang Wang, Shan Xue, Surya Nepal, and Quan Z. Sheng. 2022. eRiskCom: an e-commerce risky community detection platform. *VLDB J.* 31, 5 (2022), 1085–1101.
- [26] Yang Liu, Xiang Ao, Zidi Qin, Jianfeng Chi, Jinghua Feng, Hao Yang, and Qing He. 2021. Pick and Choose: A GNN-based Imbalanced Learning Approach for Fraud Detection. In *WWW*. 3168–3177.
- [27] Ziqi Liu, Chaochao Chen, Xinxing Yang, Jun Zhou, Xiaolong Li, and Le Song. 2018. Heterogeneous Graph Neural Networks for Malicious Account Detection. In *CIKM*. 2077–2085.
- [28] Wei Lu and Ali A. Ghorbani. 2009. Network Anomaly Detection Based on Wavelet Analysis. *EURASIP J. Adv. Signal Process.* 2009 (2009), 1–16.
- [29] Federico Monti, Davide Boscaini, Jonathan Masci, Emanuele Rodolà, Jan Svoboda, and Michael M. Bronstein. 2017. Geometric Deep Learning on Graphs and Manifolds Using Mixture Model CNNs. In *CVPR*. 5425–5434.
- [30] Mike Schuster and Kuldeep K. Paliwal. 1997. Bidirectional recurrent neural networks. *IEEE Trans. Signal Process.* 45, 11 (1997), 2673–2681.
- [31] Chuan Shi, Binbin Hu, Wayne Xin Zhao, and Philip S. Yu. 2019. Heterogeneous Information Network Embedding for Recommendation. *IEEE Trans. Knowl. Data Eng.* 31, 2 (2019), 357–370.
- [32] Kai Shu, Amy Sliva, Suhang Wang, Jiliang Tang, and Huan Liu. 2017. Fake News Detection on Social Media: A Data Mining Perspective. *SIGKDD Explor.* 19, 1 (2017), 22–36.
- [33] Jianheng Tang, Fengrui Hua, Ziqi Gao, Peilin Zhao, and Jia Li. 2023. GADBench: Revisiting and Benchmarking Supervised Graph Anomaly Detection. In *NeurIPS*.
- [34] Jianheng Tang, Jiajin Li, Ziqi Gao, and Jia Li. 2022. Rethinking Graph Neural Networks for Anomaly Detection. In *ICML*. 21076–21089.
- [35] Lloyd N Trefethen. 2019. *Approximation theory and approximation practice, extended edition*. SIAM.
- [36] Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph Attention Networks. In *ICLR*.
- [37] Binghui Wang, Neil Zhenqiang Gong, and Hao Fu. 2017. GANG: Detecting Fraudulent Users in Online Social Networks via Guilt-by-Association on Directed Graphs. In *ICDM*. 465–474.
- [38] Xiao Wang, Houye Ji, Chuan Shi, Bai Wang, Yanfang Ye, Peng Cui, and Philip S. Yu. 2019. Heterogeneous Graph Attention Network. In *WWW*. 2022–2032.
- [39] Yuchen Wang, Jinghui Zhang, Zhengjie Huang, Weibin Li, Shikun Feng, Ziheng Ma, Yu Sun, Dianhai Yu, Fang Dong, Jiahui Jin, Beilun Wang, and Junzhou Luo. 2023. Label Information Enhanced Fraud Detection against Low Homophily in Graphs. In *WWW*. 406–416.
- [40] Bingbing Xu, Huawei Shen, Qi Cao, Yunqi Qiu, and Xueqi Cheng. 2019. Graph Wavelet Neural Network. In *ICLR*.
- [41] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. 2019. How Powerful are Graph Neural Networks?. In *ICLR*.
- [42] Xiao Yu, Xiang Ren, Yizhou Sun, Quanquan Gu, Bradley Sturt, Urvashi Khandelwal, Brandon Norick, and Jiawei Han. 2014. Personalized entity recommendation: a heterogeneous information network approach. In *WSDM*. 283–292.
- [43] Chuxu Zhang, Dongjin Song, Chao Huang, Ananthram Swami, and Nitesh V. Chawla. 2019. Heterogeneous Graph Neural Network. In *KDD*. 793–803.
- [44] Yutao Zhang, Jie Tang, Zhilin Yang, Jian Pei, and Philip S. Yu. 2015. COSNET: Connecting Heterogeneous Social Networks with Local and Global Consistency. In *KDD*. 1485–1494.
- [45] Yusong Zhang, Kun Xie, Xingyi Zhang, Xiangyu Dong, and Sibow Wang. 2025. Rumor Detection on Social Media with Reinforcement Learning-based Key Propagation Graph Generator. In *WWW*. 2742–2753.
- [46] Shichao Zhu, Chuan Zhou, Shirui Pan, Xingquan Zhu, and Bin Wang. 2019. Relation Structure-Aware Heterogeneous Graph Neural Network. In *ICDM*. 1534–1539.

A Profiles of Datasets

The statistics of datasets are summarized in Table 4. For heterogeneous datasets, the #Features column in Table 4 is an array listing the feature dimension for each node type; the anomaly ratio is defined as the ratio of abnormal nodes to target type nodes, rather than the ratio to all types of nodes.

B Implementation Setup

For reproducibility, our code is available at [1]. Next, we explain how the data is split and how we set the hyper-parameters.

B.1 Data Split

The datasets are split with the open-sourced code published in GitHub repository¹ of GADBench. Settings are detailed in Table 4.

B.2 Hyper-Parameter Setting

Hyper-parameters for ChiGNN and ChiGAD experiments on homogeneous and heterogeneous graphs are summarized in Table 5 with performance evaluations in Tables 1 and 2. The best model, selected via grid search based on validation set and F1-macro, is evaluated on the test set. Baselines use default hyper-parameters suggested by their authors.

¹<https://github.com/squareRoot3/GADBench/tree/master>

Table 4: Summary of datasets.

Dataset	#Node	#Edge	#Feat.	Anomaly %	Split % (Train/Val/Test)	Relation	Feat. Type
ACM (het)	10,942	547,872	[193,115,250,95]	32.68/32.07/35.07	40/20/40	Co-authorship	Misc. Info.
R-I (het)	2,880,722	8,781,690	[14,171,10,13]	0.35	40/20/40	Account Interaction	Misc. Info.
R-II (het)	6,168,746	15,576,508	[14,171,10,13]	0.55	40/20/40	Account Interaction	Misc. Info.
Reddit	10,984	168,016	64	3.3	40/20/40	Under Same Post	Text Embedding
Tolokers	11,758	519,000	10	21.8	50/25/45	Work Collaboration	Misc. Info.
Amazon	11,944	4,398,392	25	9.5	70/10/20	Review Correlation	Misc. Info.
T-Finance	39,357	21,222,543	10	4.6	40/20/40	Transaction Record	Misc. Info.
Questions	48,921	153,540	301	3.0	50/25/25	Question Answering	Text Embedding
DGraph-Fin	3,700,550	4,300,999	17	1.3	70/15/15	Loan Guarantor	Misc. Info.
T-Social	5,781,065	73,105,508	10	3.0	40/20/40	Social Friendship	Misc. Info.

Table 5: Hyper-parameter setting of ChiGNN for experiments.

Hyper-parameter	ACM	R-I	R-II	Reddit	Tolokers	Amazon	T-Finance	Questions	DGraph-Fin	T-Social
learning rate	0.0001	0.006	0.006	0.0003	0.003	0.0003	0.003	0.0001	0.01	0.01
dropout rate	0	0	0	0	0.3	0.6	0.6	0.4	0.4	0
hidden dimension	512	25	48	64	64	64	80	400	32	20
aligned dimension	512	25	48	-	-	-	-	-	-	-
epochs	200	200	200	1700	700	1550	300	150	1200	4000
MLP layers	4	2	2	2	2	3	3	4	2	3
activation	ReLU	ReLU	ReLU	Tanh	ReLU	LeakyReLU	Tanh	ReLU	ReLU	LeakyReLU
\mathcal{F}_M	-	-	-	2-11	1-15 (odd)	1-19 (odd)	1-7 (odd)	1-15 (odd)	1-15 (odd)	1-19 (odd)
H	2.2	2.2	2.1	1.05	4	3	2.05	1.05	1.05	3.05
L	1.9	1.8	1.9	0.95	3.5	1	1.05	0.95	0.95	2.95

Table 6: Expectation and most highlighted frequency of $f_i(w)$.

i	$\mathbb{E}[f_i(w)]$	$f_i^{-1}(\max)$	i	$\mathbb{E}[f_i(w)]$	$f_i^{-1}(\max)$
1	0.6970	0.0000	16	1.5940	1.7638
2	0.9603	0.6667	32	1.7126	1.8779
4	1.2180	1.1992	64	1.7973	1.9339
8	1.4313	1.5556	128	1.8571	1.9600

C Proofs

C.1 Proof of Admissibility Condition

Since $f_i(w)$ is a real-valued function, we have $|f_i(w)|^2 = f_i(w)^2$. Thus, the integral becomes:

$$\int_0^\infty \frac{f_i(w)^2}{w} dw.$$

Substituting the expression for $f_i(w)$, we get:

$$\int_0^\infty \frac{1}{w} \left(\frac{1}{S_i} \frac{1}{2^i \Gamma(i)} (w(i+1))^{i-1} e^{-\frac{w(i+1)}{2}} \right)^2 dw.$$

Simplifying the integrand:

$$\int_0^\infty \frac{1}{w} \left(\frac{1}{S_i} \frac{1}{2^i \Gamma(i)} \right)^2 (w(i+1))^{2(i-1)} e^{-w(i+1)} dw.$$

Factoring out constants:

$$\left(\frac{1}{S_i} \frac{1}{2^i \Gamma(i)} \right)^2 (i+1)^{2(i-1)} \int_0^\infty w^{2(i-1)-1} e^{-w(i+1)} dw.$$

The integral is a Gamma function:

$$\int_0^\infty w^{2(i-1)-1} e^{-w(i+1)} dw = \frac{\Gamma(2(i-1))}{(i+1)^{2(i-1)}}.$$

Substituting this back, we obtain:

$$\left(\frac{1}{S_i} \frac{1}{2^i \Gamma(i)} \right)^2 (i+1)^{2(i-1)} \cdot \frac{\Gamma(2(i-1))}{(i+1)^{2(i-1)}} = \left(\frac{1}{S_i} \frac{1}{2^i \Gamma(i)} \right)^2 \Gamma(2(i-1)).$$

Thus, we have

$$\int_0^\infty \frac{|f_i(w)|^2}{w} dw = \left(\frac{1}{S_i} \frac{1}{2^i \Gamma(i)} \right)^2 \Gamma(2(i-1)) < \infty.$$

So the Admissibility Condition is satisfied. \square

C.2 Empirical Proof of Chi-Square filters' Broad Coverage on Diverse Bands

Empirical results in Table 6 show that the focus of $\{f_i\}$ broadly covers diverse bands over the full range of $[0, 2]$.

C.3 Proof of Theorem 1

We first prove the case of aligning 2 graph signals, and then naturally extend it to the case of w signals. Supposing the input signals $\mathbf{x}_1, \mathbf{x}_2$ has been normalized, i.e. $\mathbf{x}_1^T \mathbf{x}_1 = \mathbf{x}_2^T \mathbf{x}_2 = 1$, the aligned feature \mathbf{x}_a is computed by the form of weighted sum as:

$$\mathbf{x}_a = w_1 \mathbf{x}_1 + w_2 \mathbf{x}_2.$$

Let the high-frequency area of $\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_a$ be simply S_1, S_2 and S_a . We have

$$S_a = \frac{\mathbf{x}_a^T \mathbf{L} \mathbf{x}_a}{\mathbf{x}_a^T \mathbf{x}_a} = \frac{w_1^2 S_1 + w_2^2 S_2 + 2w_1 w_2 \mathbf{x}_1^T \mathbf{L} \mathbf{x}_2}{w_1^2 + w_2^2 + 2w_1 w_2 \mathbf{x}_1^T \mathbf{x}_2}. \quad (13)$$

Following the existing studies [3, 21, 34], the graph spectral decomposition of the Laplacian matrix L can be noted as

$$L = \mathbf{U} \mathbf{\Lambda} \mathbf{U}^T.$$

Table 7: Experimental results of ablation study on the ACM dataset.

Model	AUROC	AUPRC	F1-macro	Recall
ChiGAD	0.9702	0.9518	0.9195	0.8947
w/o Multi-Graph Chi-Square Filter	0.9538	0.9236	0.8889	0.8499
w/o Interactive Meta-Graph Chi-Square Convolution	0.8969	0.8410	0.7988	0.7343
w/o Contribution-Informed Cross-Entropy Loss	0.9237	0.8449	0.8526	0.8070
Variant of ChiGAD (Beta filter)	0.9411	0.9093	0.8735	0.8371

Since L is a semi-positive definite matrix, the diagonal elements of Λ are all non-negative, and thereby so is $\Lambda^{\frac{1}{2}}$. Based on this finding, it can be deduced that

$$L^{\frac{1}{2}} = U\Lambda^{\frac{1}{2}}U^T$$

is also semi-positive definite. According to the Cauchy-Schwartz Theorem, it can be deduced that:

$$\begin{aligned} (\mathbf{x}_1^T L \mathbf{x}_2)^2 &= [(L^{\frac{1}{2}} \cdot \mathbf{x}_1)^T (L^{\frac{1}{2}} \cdot \mathbf{x}_2)]^2 \\ &\leq \sqrt{[(L^{\frac{1}{2}} \cdot \mathbf{x}_1)^T (L^{\frac{1}{2}} \cdot \mathbf{x}_1)]} \cdot \sqrt{[(L^{\frac{1}{2}} \cdot \mathbf{x}_2)^T (L^{\frac{1}{2}} \cdot \mathbf{x}_2)]} \\ &= |\mathbf{x}_1^T L \mathbf{x}_1| \cdot |\mathbf{x}_2^T L \mathbf{x}_2| = S_1 S_2, \end{aligned}$$

and thereby

$$\mathbf{x}_1^T L \mathbf{x}_2 \in [-\sqrt{S_1 S_2}, \sqrt{S_1 S_2}]. \quad (14)$$

Similarly, it is easy to see that:

$$(\mathbf{x}_1^T \mathbf{x}_2)^2 \leq |\mathbf{x}_1^T \mathbf{x}_1| |\mathbf{x}_2^T \mathbf{x}_2| = 1, \quad (15)$$

and thereby

$$\mathbf{x}_1^T \mathbf{x}_2 \in [-1, 1].$$

Combining Equations 13, 14 and 15, we can derive a lower bound of S_a in the following two cases:

Case 1: $w_1 \cdot w_2 \geq 0$. In this case, it can be inferred that

$$\begin{aligned} S_a \geq \mathcal{LB}_{c-1} &:= \frac{w_1^2 S_1 + w_2^2 S_2 - 2w_1 w_2 \sqrt{S_1 S_2}}{w_1^2 + w_2^2 + 2w_1 w_2} \\ &= \left(\frac{w_1}{w_1 + w_2} \sqrt{S_1} - \frac{w_2}{w_1 + w_2} \sqrt{S_2} \right)^2; \end{aligned}$$

Case 2: $w_1 \cdot w_2 < 0$. In this case, it can be inferred that

$$\begin{aligned} S_a \geq \mathcal{LB}_{c-2} &:= \frac{w_1^2 S_1 + w_2^2 S_2 + 2w_1 w_2 \sqrt{S_1 S_2}}{w_1^2 + w_2^2 - 2w_1 w_2} \\ &= \left(\frac{w_1}{w_1 - w_2} \sqrt{S_1} + \frac{w_2}{w_1 - w_2} \sqrt{S_2} \right)^2. \end{aligned}$$

By symmetry, it can be obtained that $\mathcal{LB}_{c-1} = \mathcal{LB}_{c-2}$. Hence we only consider **Case 1**, where $w_1, w_2 \geq 0$. Let $t = \frac{w_1}{w_2} \geq 0$, then

$$f(t) := \mathcal{LB}_{c-1} = \left(\frac{t\sqrt{S_1} - \sqrt{S_2}}{t+1} \right)^2,$$

and

$$f'(t) = 2 \left(\frac{t\sqrt{S_1} - \sqrt{S_2}}{t+1} \right) \frac{\sqrt{S_1} + \sqrt{S_2}}{(t+1)^2}. \quad (16)$$

Equation 16 reveals that the lower bound $f(t)$ monotonically decreases on $[0, \sqrt{\frac{S_2}{S_1}}]$ where $f'(t) \leq 0$ while monotonically increases

on $(\sqrt{\frac{S_2}{S_1}}, +\infty)$ where $f'(t) > 0$. With these findings, it is clear that

$$f(0) = S_2, \quad \lim_{t \rightarrow +\infty} f(t) = S_1$$

serve as the potential upper bound of $f(t)$, and notably, they are both achievable or arbitrarily approximated by adjusting t , i.e. the weights w_1, w_2 . Specifically, if $S_1 \leq S_2$, the adjustment of $w_2 \gg w_1 \geq 0$ deserves to be implemented, and vice versa ($w_1 \gg w_2 \geq 0$) if $S_1 > S_2$. To this end, the lower bound $f(t)$ of S_a is sufficiently close to or even reaches an ideal upper bound $\max\{S_1, S_2\}$, i.e.,

$$S_{high}(\mathbf{x}_a, L) = \max \{S_{high}(\mathbf{x}_i, L)\}_{i=1}^2$$

Therefore, the proof for aligning two graph signals has been completed. Then, we extend this to $k \in \mathbb{Z}^+$ graph signals. Assuming the theorem holds with $k-1$ graph signals:

$$S_{high}(\mathbf{x}_a, L) = \max \{S_{high}(\mathbf{x}_i, L)\}_{i=1}^{k-1}.$$

Assume the following weighted sum $\mathbf{x}_a = \sum_{i=1}^{k-1} w_i^* \mathbf{x}_i$ satisfies

$$S_{high}(\mathbf{x}_a, L) = \max \{S_{high}(\mathbf{x}_i, L)\}_{i=1}^{k-1}.$$

Considering \mathbf{x}_a as a new graph signal, the alignment of \mathbf{x}_a and \mathbf{x}_k degrades to the demonstrated case where there are only two graph signals. The aligned result of $\mathbf{x}_a, \mathbf{x}_k, \mathbf{x}'_a$ can be inferred that:

$$\begin{aligned} S_{high}(\mathbf{x}'_a, L) &= \max \{S_{high}(\mathbf{x}_a, L), S_{high}(\mathbf{x}_k, L)\} \\ &= \max \{ \max \{S_{high}(\mathbf{x}_i, L)\}_{i=1}^{k-1}, S_{high}(\mathbf{x}_k, L) \} \\ &= \max \{S_{high}(\mathbf{x}_i, L)\}_{i=1}^k. \end{aligned}$$

Since \mathbf{x}'_a acts as a linear alignment of all the k graph signals $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_k$, Theorem 1 is fully demonstrated. \square

D Experimental Results of Ablation Study

Ablation study for ChiGAD is shown in Table 7. We can observe that ChiGAD outperforms its ablated variants on all metrics (AUROC, AUPRC, F1-macro, and Recall), indicating that each of its components contributes to improved anomaly detection. Notably, removing the Multi-Graph Chi-Square Filter or the Interactive Meta-Graph Convolution leads to the most substantial drop in performance, underscoring their critical role in capturing heterogeneous information and aligning feature dimensions across node types.

Excluding the Contribution-Informed Cross-Entropy Loss diminishes Recall, which shows that weighting anomalies based on their frequency contributions helps the model focus on challenging yet vital abnormal nodes. Furthermore, ChiGAD outperforms a Beta filter variant, suggesting that the Chi-Square filter's unique properties are effective in identifying anomalies.