

# A Label-Free Heterophily-Guided Approach for Unsupervised Graph Fraud Detection

Junjun Pan<sup>1</sup>, Yixin Liu<sup>1</sup>, Xin Zheng<sup>1</sup>, Yizhen Zheng<sup>2</sup>,  
Alan Wee-Chung Liew<sup>1</sup>, Fuyi Li<sup>3</sup>, Shirui Pan<sup>1\*</sup>

<sup>1</sup>School of Information and Communication Technology, Griffith University, Queensland, Australia

<sup>2</sup>Faculty of Information Technology, Monash University, Melbourne, Australia

<sup>3</sup>Faculty of Health and Medical Sciences, The University of Adelaide, South Australia, Australia

{junjun.pan, yixin.liu, xin.zheng, a.liew, s.pan}@griffith.edu.au, yizhen.zheng1@monash.edu, fuyi.li@adelaide.edu.au

## Abstract

Graph fraud detection (GFD) has rapidly advanced in protecting online services by identifying malicious fraudsters. Recent supervised GFD research highlights that heterophilic connections between fraudsters and users can greatly impact detection performance, since fraudsters tend to camouflage themselves by building more connections to benign users. Despite the promising performance of supervised GFD methods, the reliance on labels limits their applications to unsupervised scenarios; Additionally, accurately capturing complex and diverse heterophily patterns without labels poses a further challenge. To fill the gap, we propose a **Heterophily-guided Unsupervised Graph fraud dEtection** approach (HUGE) for unsupervised GFD, which contains two essential components: a heterophily estimation module and an alignment-based fraud detection module. In the heterophily estimation module, we design a novel label-free heterophily metric called HALO, which captures the critical graph properties for GFD, enabling its outstanding ability to estimate heterophily from node attributes. In the alignment-based fraud detection module, we develop a joint MLP-GNN architecture with ranking loss and asymmetric alignment loss. The ranking loss aligns the predicted fraud score with the relative order of HALO, providing an extra robustness guarantee by comparing heterophily among non-adjacent nodes. Moreover, the asymmetric alignment loss effectively utilizes structural information while alleviating the feature-smooth effects of GNNs. Extensive experiments on 6 datasets demonstrate that HUGE significantly outperforms competitors, showcasing its effectiveness and robustness.

## Introduction

With the advancement of information technology, graph-structured data have become ubiquitous in online services such as e-commerce (Motie and Raahemi 2023; Zheng et al. 2022b, 2024b, 2022c; Wang et al. 2024a; Zheng et al. 2024a) and social media (Deng et al. 2022). However, this widespread usage has led to a significant increase in various malicious activities, including fraud, spam, and fake reviews, making detecting such activities crucial. To address this issue, graph fraud detection (GFD) is an emerging research topic that aims to identify these suspicious activities.

\*Corresponding Author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

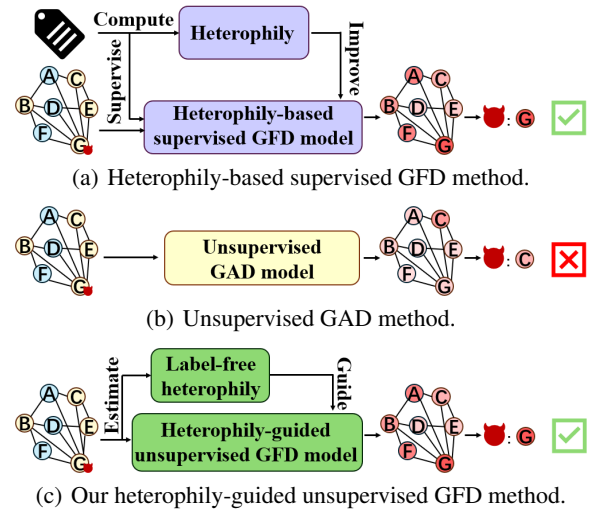


Figure 1: Concept maps of existing GFD/GAD methods and our proposed unsupervised GFD method HUGE.

By analyzing both structural patterns and attribute contexts within networks, GFD ensures the integrity of online applications (Motie and Raahemi 2023; Deng et al. 2022; Yu et al. 2022).

Recently, supervised GFD methods based on graph neural networks (GNNs) have shown significant progress and attracted considerable research interest (Dou et al. 2020; Gao et al. 2023a; Li et al. 2022). One of the mainstream research focuses is on addressing the graph **heterophily** introduced by fraudsters (Gao et al. 2023a; Zheng et al. 2022a), i.e., fraudsters tend to establish connections with benign users to avoid being detected for concealing their suspicious activities. In this way, conventional GNN models struggle to differentiate between the representations of fraudsters and benign users due to the smooth effect of message passing (Cai and Wang 2020), resulting in sub-optimal GFD performance. In light of this, various strategies have been proposed to mitigate the impact of heterophily, such as dropping edges (Gao et al. 2023a), or using spectral GNNs (Tang et al. 2022). Despite the promising performance of these methods, they rely heavily on annotated labels of both fraudsters and benign users for modeling the heterophily property. How-

ever, in real-world scenarios, obtaining annotated fraudsters typically requires significant human effort, making supervised GFD methods impractical when labels are unavailable. Therefore, developing *unsupervised GFD methods* to avoid label reliance can be a promising research direction.

Recently, several studies have attempted to address the challenge of unsupervised GFD under graph heterophily by leveraging graph anomaly detection (GAD). GAD can be regarded as a superset of GFD, as it not only seeks to identify fraudsters but also any rare events or patterns within graph data (Ma et al. 2021), thereby making it a viable solution for unsupervised GFD. However, these methods have primarily demonstrated promising results on artificial datasets within the general GAD context. The artificial anomalous patterns (e.g., dense subgraph and feature outliers) are relatively straightforward to identify (Ding et al. 2019; Gu and Zou 2024). In contrast, in real-world fraud detection, fraudsters often introduce more heterophilic connections with benign users to conceal their activities. To address these challenges, recent works (He et al. 2024; Qiao and Pang 2024) have sought to iteratively reduce heterophily from graph structure by estimating pseudo-labels. Nevertheless, the iterative refinement and multiple-round pseudo-label computation limit their time and memory efficiency. Moreover, estimating pseudo-labels using GNNs may still suffer from feature-smoothing effects, potentially misleading the subsequent iterations. The general comparison between existing solutions and our proposed method is demonstrated in Figure 1.

In light of the insufficiency of existing solutions for the unsupervised GFD problem under graph heterophily, a natural research question arises:

*Can we explicitly measure heterophily without annotated labels and utilize it to guide the learning of the GFD model?*

To answer this question, we identify two challenging aspects for effective unsupervised GFD:

**Challenge 1 - Precise unsupervised heterophily estimation:** the calculation of heterophily often relies on labels (Zheng et al. 2022a) that are unavailable in unsupervised GFD scenarios. While several attribute-level distance metrics have the potential to indicate heterophily, they hardly satisfy some crucial properties for effective GFD (e.g. Euclidean distance is unbounded, cosine distance is scale insensitive, etc.), limiting their applicability in unsupervised GFD.

**Challenge 2 - Enabling effective heterophily guidance in GFD:** even though a proper unsupervised heterophily metric can be established, how to effectively use them for subsequent GAD remains a challenging task. Since the heterophily is estimated from attributes, it provides not only clues about fraudsters but also noise. Therefore, simply following previous supervised GFD methods may lead to sub-optimal results for unsupervised GFD.

To address the above challenges, in this paper, we proposed a novel **Heterophily-guided Unsupervised Graph fraud dEtection** method (HUGE for short), which contains two essential components: a heterophily estimation module and an alignment-based fraud detection module, which provides a powerful solution for unsupervised GFD that overcomes previous challenges. Specifically, we analyze the desired properties that label-free heterophily measures should

possess for GFD. Based on these, we design a novel metric called HALO to address **Challenge 1**, which captures the critical graph properties for GFD, enabling its outstanding ability to estimate heterophily from node attributes. In the alignment-based fraud detection module, we develop a joint MLP-GNN architecture with ranking loss and asymmetric alignment loss to address **Challenge 2**. The ranking loss aligns the order of predicted fraud scores with that of HALO to guide the model training process. Utilizing the rank of heterophily as guidance not only improves the robustness by mitigating irrelevant information, but also introduces extra global heterophily constraints by comparing heterophily among non-adjacent nodes. Moreover, the asymmetric alignment loss effectively extracts structural information from GNN to MLP by aligning the neighbor inconsistency distributions, alleviating the feature-smooth effect introduced by message passing. For evaluation, we conduct extensive experiments to demonstrate the superiority of HUGE over state-of-the-art methods on six real-world GFD datasets. In summary, our contributions are three-fold:

- We derive a new label-free heterophily metric termed HALO to accurately estimate heterophily by capturing the critical graph properties for unsupervised GFD.
- We develop a novel unsupervised GFD framework, **Heterophily-guided Unsupervised Graph fraud dEtection** (HUGE). We first estimate heterophily from attributes with HALO, then utilize a joint MLP-GNN architecture with ranking loss and asymmetric alignment loss to guarantee robustness and effectiveness. Finally, the local inconsistency scores of MLP embeddings are used as fraud scores to spot fraudsters.
- We conduct extensive experiments to demonstrate the superior performance of HUGE over the state-of-the-art methods on six real-world GFD datasets under unsupervised learning scenarios.

## Related Work

**Graph Fraud Detection (GFD)** aims to identify fraudulent activities from graph-structured real-world systems, including financial fraud (Motie and Raahemi 2023), spamming (Deng et al. 2022), and fake reviews (Yu et al. 2022). Various techniques have been utilized to detect the fraudsters, including attention (Liu et al. 2020), sampling (Liu et al. 2021a,a) and multi-view learning (Zhong et al. 2020). Although these methods have achieved promising results, they suffer from heterophily-caused problem (Dou et al. 2020), i.e., fraudster nodes tend to build heterophilic connections with benign user nodes to make them indistinguishable from the majority. Recent studies have developed many strategies to mitigate the negative impact of heterophily with node labels. For example, GHRN (Gao et al. 2023a) reduce heterophily by pruning the graph, while BWGNN (Tang et al. 2022) leverage spectral GNNs to better capture high-frequency features associated with heterophily. GDN (Gao et al. 2023b), GAGA (Wang et al. 2023), and PMP (Zhuo et al. 2024) overcome the feature-smoothing effect by crafting advanced encoders that sharpen feature separation, while ConsisGAD (Chen et al. 2024b) directly utilizes annotated

Metric	Definition	Boundedness	Minimal Agreement	Monotonicity	Equal Attribute Tolerance
Euclidean Distance	$\ \mathbf{x}_i - \mathbf{x}_j\ $	✗	✓	✓	✓
Cosine Distance	$-\frac{\mathbf{x}_i^T \cdot \mathbf{x}_j}{\ \mathbf{x}_i\  \ \mathbf{x}_j\ }$	✓	✗	✗	✗
Attr. Het. Rate (Yang et al. 2021)	$\frac{\sum_{k=1}^d \mathbf{1}[x_{i,k} \neq x_{j,k}]}{d}$	✓	✓	✗	✗
<b>HALO (ours)</b>	Eq. (2)	✓	✓	✓*	✓

Table 1: Comparison of unsupervised edge-level heterophily metrics. ✓\* denotes that the metric satisfies the property under relaxed constraints. Vectors  $\mathbf{x}_i, \mathbf{x}_j$  denote attributes of nodes  $v_i, v_j$  respectively. Attr. Het. Rate refers to attribute heterophily rate. The indicator function  $\mathbf{1}[x_{i,k} \neq x_{j,k}]$  returns 1 when the condition  $x_{i,k} \neq x_{j,k}$  is satisfied, and 0 otherwise.

labels to create more effective graph augmentation method. Despite their promising results, the reliance on labels in existing methods restricts their applicability in unsupervised scenarios. Therefore, in this work, we aim to address the pressing need for developing unsupervised GFD methods.

**Graph Anomaly Detection (GAD)** is a broader concept than GFD, aiming to identify not only fraudsters but also any rare and unusual patterns that significantly deviate from the majority in graph data (Ding et al. 2019; Zheng et al. 2021b; Wang et al. 2024b; Liu et al. 2024b; Cai et al. 2024; Cai and Fan 2022; Zhang et al. 2024; Liu et al. 2023). Therefore, GAD techniques can be directly applied to GFD, especially in unsupervised learning scenarios (Li et al. 2024; Liu et al. 2024a). Given the broad scope of GAD and the difficulty in obtaining real-world anomalies, many unsupervised GAD methods have been designed and evaluated on several datasets with artificially injected anomalies (Ding et al. 2019; Liu et al. 2021b; Jin et al. 2021; Zheng et al. 2021a; Duan et al. 2023a,b; Pan et al. 2023). Despite their decent performances, these methods rely on the strong homophily of the datasets with injected anomalies, which limits their applications under graph heterophily (Zheng et al. 2022a, 2023). Recent studies explore this issue and suggest using estimated anomaly scores as pseudo-labels to mitigate the negative impacts of heterophily, e.g., dropping edges (He et al. 2024; Qiao and Pang 2024) and adjusted message passing (Chen et al. 2024a). However, the estimated anomaly score might not be an effective indicator for modeling the heterophily under the homophily-based message-passing GNNs (Zhu et al. 2022), leading to suboptimal GFD performance (Zhao and Akoglu 2020). Hence, we introduce a novel heterophily measurement with desired properties as effective guidance for a powerful unsupervised GFD model.

More detailed related works are available in Appendix A.

## Preliminary

**Notations.** Let  $G = (\mathcal{V}, \mathcal{E}, \mathbf{X})$  represents an attributed graph, where  $\mathcal{V}$  is the set of nodes and  $\mathcal{E}$  is the set of edges. We denote the number of nodes and edges as  $n$  and  $m$ , respectively.  $\mathbf{X} \in \mathcal{R}^{n \times d}$  is the attribute matrix, with  $d$  being the dimension of the attributes. The attribute vector of the  $i$ -th node  $v_i$  is represented by its  $i$ -th row, i.e.,  $\mathbf{x}_i$ . The graph structure can also be represented as adjacent matrix  $\mathbf{A} \in \mathcal{R}^{n \times n}$ , with  $\mathbf{A}_{i,j} = 1$  if there exists an edge between nodes  $v_i$  and  $v_j$ , and  $\mathbf{A}_{i,j} = 0$  otherwise. We denote the set of neighbors of the  $i$ -th node as  $\mathcal{N}(v_i) = \{v_j | \mathbf{A}_{i,j} = 1\}$ .

**Problem Definition.** In the context of graph fraud detection (GFD), each node  $v_i \in \mathcal{V}$  has a label  $y_i \in \{0, 1\}$ , where  $y_i = 0$  indicates a benign node and  $y_i = 1$  represents a fraudulent node. An assumption is that the number of benign nodes is significantly greater than the number of fraudulent nodes. In unsupervised GFD scenarios, labels are unavailable during the training stage. The goal of a GFD method is to identify whether a node  $v_i$  is suspicious by predicting a fraudulent score  $s_i^{\text{fraud}}$ , where a higher score indicates a greater likelihood that the node is fraudulent.

## Methodology

In this section, we introduce the proposed unsupervised graph fraud detection (GFD) approach termed **Heterophily-guided Unsupervised Graph fraud dEtection (HUGE)**. As illustrated in Figure 2, HUGE consists of two components at the highest level: label-free heterophily estimation module and alignment-based fraud detection module. We first estimate the heterophily using a novel label-free heterophily measure termed HALO. Then, the fraud detection model is trained following the guidance of the relative order of HALO with a ranking loss and an asymmetric alignment loss, leading to an effective unsupervised GFD framework. More details of the proposed HUGE are presented as follows.

### Label-free Heterophily Estimation

GFD contends with a well-known challenge named the camouflage effects (Dou et al. 2020), wherein fraudsters tend to connect with normal nodes to hide their suspicious features. To address this issue, heterophilic edges serve as a crucial hint, revealing the feature of connections that fraudsters attempt to establish. However, under unsupervised application scenarios, it becomes significantly challenging to directly compute heterophily from attributes. Hence, in this work, we proposed to comprehensively and precisely estimate heterophily from node attributes.

In this subsection, we first discuss a set of universally desired properties for the estimation of heterophily in unsupervised GFD and provide a detailed review of the aforementioned insufficiency. We then introduce a new label-free heterophily measure named HALO, which satisfies these properties and provides better guidance for detecting camouflaged fraudsters.

**Desirable properties for heterophily measures.** Measuring the inconsistency between two nodes is the fundamental building block for modeling heterophily. While defin-

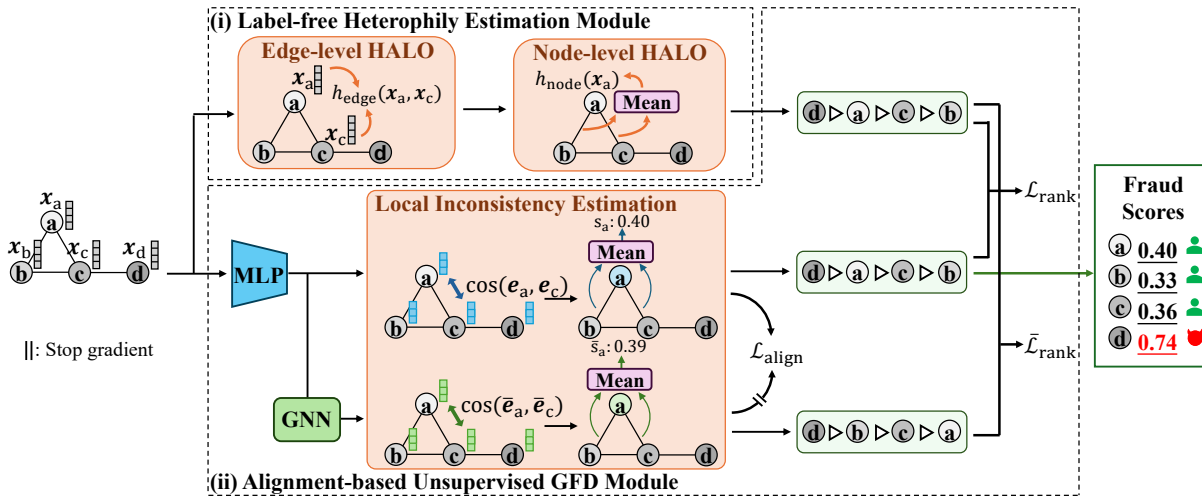


Figure 2: Workflow of HUGE. Our *label-free heterophily estimation module* estimates node heterophily using attributes. Then, in the *alignment-based unsupervised GFD module*, a joint MLP-GNN architecture is trained through ranking and asymmetric alignment losses. The ranking loss ensures the predicted inconsistency score aligns with the heterophily order, while the asymmetric alignment loss matches the neighbor inconsistency distribution of the MLP encoder to that of the GNN encoder. In evaluation phase, the local inconsistency scores generated by the MLP encoder are used as the final fraud scores.

ing heterophily (Zheng et al. 2022a) from annotated labels is straightforward, extending its definition using node attributes is challenging, as attributes are continuous variables with multiple dimensions. Moreover, fraudsters may intentionally modify their attributes to appear normal (Dou et al. 2020), complicating heterophily estimation. Motivated by (Platonov et al. 2024), we summarize a set of desirable properties for edge-level label-free heterophily  $h_{\text{edge}}$ .

- **Boundedness.** This property requires the measurement to have constant upper bounds  $c_{\max}$  and lower bounds  $c_{\min}$  for any node attributes. Formally, we have  $c_{\min} \leq h_{\text{edge}}(\mathbf{x}_i, \mathbf{x}_j) < c_{\max}, \forall \mathbf{x}_i, \mathbf{x}_j \in \mathcal{R}^d$ . Unbounded measurements may be overly sensitive to a few attributes with extreme values. For example, overemphasizing the frequency of sending messages may misidentify highly engaged users as spammers.
- **Minimal Agreement.** In GFD, if two users share identical profiles, the heterophily between them should reach its minimum value without any other information. This property is defined as minimal agreement, i.e.,  $c_{\min} = h_{\text{edge}}(\mathbf{x}_i, \mathbf{x}_j) \iff \mathbf{x}_i = \mathbf{x}_j$ .
- **Monotonicity.** To precisely distinguish heterophilic connections from normal edges, a promising heterophily metric should increase as the distance between attributes increases. While it is challenging to model it without prior knowledge of attribute distribution, we provide a definition that focuses on local inconsistency:

$$\forall \mathbf{x}_i \neq \mathbf{x}_j, 1 \leq k \leq d, \text{ we have:}$$

$$\text{If } x_{i,k} > x_{j,k} \text{ then: } \frac{\partial h_{\text{edge}}(\mathbf{x}_i, \mathbf{x}_j)}{\partial x_{i,k}} > 0 \quad (1)$$

$$\text{Otherwise: } \frac{\partial h_{\text{edge}}(\mathbf{x}_i, \mathbf{x}_j)}{\partial x_{i,k}} < 0.$$

- **Equal Attribute Tolerance.** Fraudsters often mimic the attributes of benign users to evade detection. To avoid bias from globally shared attributes, similar to empty class tolerance (Platonov et al. 2024), we define equal attribute tolerance as:  $\forall k, h_{\text{edge}}(\mathbf{x}_i, \mathbf{x}_j) = h_{\text{edge}}([\mathbf{x}_i^T \parallel k]^T, [\mathbf{x}_j^T \parallel k]^T)$ , where  $\parallel$  denotes the concatenate operation.

Based on the aforementioned properties, we analyze the existing unsupervised metrics and summarize their properties in Table 1. Detailed proof and examples are provided in Appendix B. While these metrics can estimate heterophily to some extent, they fail to capture important clues essential for GFD. For example, cosine distance and attribute heterophily rate may overlook scale when attributes are sparse, diminishing their effectiveness in detecting accounts with numerous fake followers. Conversely, Euclidean distance may overemphasize attributes with strong disparity, such as the number of posts or likes, thereby neglecting the semantic differences in posts. Such insufficiencies limit their ability to estimate heterophily in GFD tasks, particularly when fraudsters attempt to conceal themselves among benign users.

**Harmonic Label-free Heterophily.** With regard to the limitations of existing measures, an ideal heterophily metric for GFD should be robust across diverse attribute distributions in various datasets, while remaining sensitive to suspicious attributes to effectively identify fraudulent nodes. To bridge the gap, we propose an attribute-based heterophily metric termed **H**Armonic **L**abel-free heter**O**phily (HALO), which can be written by:

$$\text{HALO}(\mathbf{x}_i, \mathbf{x}_j) = \frac{\|\hat{\mathbf{x}}_i - \hat{\mathbf{x}}_j\|_2}{(\|\hat{\mathbf{x}}_i\|_2^2 + \|\hat{\mathbf{x}}_j\|_2^2 + \epsilon)^{\frac{1}{2}}}, \quad (2)$$

where  $\hat{\mathbf{x}}_i = \text{abs}(\mathbf{x}_i - \mathbf{x}_j) \odot \mathbf{x}_i$  and  $\hat{\mathbf{x}}_j = \text{abs}(\mathbf{x}_i - \mathbf{x}_j) \odot \mathbf{x}_j$ . Here,  $\text{abs}(\cdot)$  and  $\odot$  denote element-wise absolute value and

multiplication, respectively, and  $\epsilon$  is a small positive number to avoid division by zero. To calculate HALO, we first rescale attributes based on their pairwise absolute distance, and then compute the normalized Euclidean distance between the rescaled attributes to ensure HALO is bounded.

With theoretical analysis, we prove that HALO strictly satisfies boundedness, minimal agreement and equal attributes tolerance, and satisfies monotonicity under relaxed constraints. The detailed proofs are included in Appendix C. Satisfying boundedness ensures HALO will not overemphasize a few attributes, while minimal agreement and equal attributes tolerance ensure the metric remains well-defined and robust across various GFD datasets with different attribute distributions. We relax strict monotonicity to enhance sensitivity to attributes that may camouflage fraudsters, aligning with our objective of using heterophily to guide unsupervised GFD.

To compute node-level heterophily, we simply average the edge-level HALO between a node and its neighbors:

$$h_i = \text{HALO}(v_i) = \frac{\sum_{v_j \in \mathcal{N}(v_i)} \text{HALO}(\mathbf{x}_i, \mathbf{x}_j)}{|\mathcal{N}(v_i)|}, \quad (3)$$

which provides a straightforward and intuitive measure of local heterophily within the graph. While it is possible to design more theoretically grounded node-level heterophily measures (Platonov et al. 2024), we leave this as a direction for future research.

### Alignment-based Graph Fraud Detection

To effectively utilize HALO to enhance the performance of unsupervised GFD, we propose an alignment-based graph fraud detection module, which consists of a joint MLP-GNN architecture with ranking loss and asymmetric alignment loss. Specifically, we first learn node embeddings from the MLP encoder to estimate local inconsistency scores as indicators of potential fraudsters. Subsequently, we align the order of local inconsistency scores to the order of computed heterophily HALO through the ranking loss, which reduces the influences of noise. Next, we use the GNN encoder to capture the graph structural information, and design the asymmetric alignment loss for aligning the neighbor inconsistency distribution between the MLP and the GNN encoders. The asymmetric alignment loss not only effectively integrates structure information into the MLP encoder, but also alleviates the feature-smoothing effects of GNN.

**Local Inconsistency Estimation.** In the first step, we obtain the node embedding using an MLP encoder, which can be written as  $\mathbf{E} = \text{MLP}(\mathbf{X})$ , where  $\mathbf{E} = \{\mathbf{e}_i\}_{i=1}^n$  and  $\mathbf{e}_i$  represents the embedding of node  $v_i$ . Utilizing an MLP as an encoder ensures that the embeddings maintain their distinctiveness, thereby preventing the discriminative indicators of fraudsters from being smoothed out by message passing. After that, we acquire the local inconsistency score  $s_i$  for node  $v_i$  by calculating the average cosine similarity between the node and its neighbors:

$$s_i = -\frac{\sum_{v_j \in \mathcal{N}(v_i)} \cos(\mathbf{e}_i, \mathbf{e}_j)}{|\mathcal{N}(v_i)|}, \quad (4)$$

where  $\cos(\cdot, \cdot)$  denotes the cosine similarity. This local inconsistency score effectively distinguishes suspicious attributes from benign ones, enabling us to effectively identify fraudsters. Therefore, we use it as the final anomaly score, i.e.,  $s_i^{\text{fraud}} = s_i$ , ensuring effective unsupervised GFD.

**Heterophily-guided Ranking Loss.** While HALO provides useful clues for GFD, directly utilizing the heterophily as the fraud score or regression target may yield sub-optimal results. As the heterophily is estimated from attributes, it often contains redundant information, such as noise or variation on attributes that are irrelevant to GFD. Such unreliable information may introduce noisy signals and reduce overall performance (Burgess et al. 2005). Moreover, heterophily only captures local inconsistencies, whereas global information is also essential for detecting the most malicious entities for unsupervised GFD problem (Jin et al. 2021). To address this, we incorporate concepts from information retrieval (Burgess et al. 2005) and employ a ranking loss to align the scores predicted by the learnable GFD model with the order of estimated heterophily:

$$\mathcal{L}_{\text{rank}}^+ = \frac{1}{|\mathcal{B}|^2} \sum_{v_i \in \mathcal{B}} \sum_{v_j \in \mathcal{B}} (1[h_j > h_i](s_i - s_j) + \log(1 + e^{-(s_i - s_j)})), \quad (5)$$

where  $\mathcal{B}$  represents a batch of sampled nodes, and  $1[\cdot]$  is the indicator function which returns 1 when the condition is satisfied and 0 otherwise. By avoiding direct mapping the scores to specific values, the ranking loss ensures the robustness of the GFD model and mitigates the influence of irrelevant information and bias (Burgess et al. 2005). Furthermore, the pairwise ordering relation allows comparisons between non-adjacent nodes, providing valuable global insights and enhancing the detection of fraudulent entities.

In addition to the order of heterophily, we highlight another implicit but important prior inspired by the homophily assumption (Zheng et al. 2022a): the local inconsistency score  $s_i$  between a node and its neighbor should be smaller than the inconsistency score between a node and its non-neighbor  $s_i^-$ . Therefore, by integrating this hidden constraint into the ranking loss, we can obtain:

$$\mathcal{L}_{\text{rank}}^- = \frac{1}{|\mathcal{B}|} \sum_{v_i \in \mathcal{B}} (s_i - s_i^- + \log(1 + e^{-(s_i - s_i^-)})), \quad (6)$$

$$s_i^- = -\frac{\sum_{v_j \in \mathcal{B} \setminus \mathcal{N}(v_i)} \cos(\mathbf{e}_i, \mathbf{e}_j)}{|\mathcal{B} \setminus \mathcal{N}(v_i)|}, \quad (7)$$

where  $\cos(\cdot, \cdot)$  denotes the cosine similarity. Finally, we use the average of  $\mathcal{L}_{\text{rank}}^+$  and  $\mathcal{L}_{\text{rank}}^-$  as the final rank loss:

$$\mathcal{L}_{\text{rank}} = \frac{\mathcal{L}_{\text{rank}}^+ + \mathcal{L}_{\text{rank}}^-}{2}. \quad (8)$$

**Asymmetric Alignment Loss.** Apart from attributes, the graph structure offers valuable information for unsupervised GFD. To leverage this without obscuring fraud clues through message passing, we introduce an auxiliary GNN after the MLP encoder. We then employ an asymmetric alignment loss to extract structural knowledge from GNN to MLP by aligning their neighbor inconsistency distribution (NID).

Methods	Amazon		Facebook		Reddit		YelpChi		AmazonFull		YelpChiFull	
	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC	AUROC	AUPRC
DOMINANT	0.4845	0.0589	0.4372	0.0194	0.5588	0.0371	0.3993	0.0385	0.4496	0.0574	OOM	
CoLA	0.4734	0.0683	0.8366	0.2155	<b>0.6032</b>	0.0440	0.4336	0.0448	0.2110	0.0536	<u>0.4911</u>	<u>0.1419</u>
ANEMONE	0.5757	0.1054	0.8385	0.2320	0.5853	0.0417	0.4432	0.0473	0.6056	<u>0.2396</u>	0.4601	0.1305
GRADATE	0.5539	0.0892	0.8809	0.3560	0.5671	0.0389	OOM		OOM		OOM	
ADA-GAD	0.5240	0.1108	0.0756	0.0122	0.5610	0.0382	OOM		OOM		OOM	
GADAM	0.6167	0.0857	0.9539	0.3630	0.5809	0.0465	0.4177	0.0423	0.4457	0.0566	0.4797	0.1351
TAM	<u>0.7126</u>	<u>0.2915</u>	0.8895	0.1937	0.5748	0.0438	<u>0.5473</u>	<b>0.0780</b>	<u>0.6442</u>	0.2188	OOM	
<b>HUGE (ours)</b>	<b>0.8516</b>	<b>0.6672</b>	<b>0.9760</b>	<b>0.3674</b>	<u>0.5906</u>	<b>0.0511</b>	<b>0.6013</b>	<u>0.0708</u>	<b>0.8892</b>	<b>0.7668</b>	<b>0.5767</b>	<b>0.1869</b>

Table 2: Unsupervised GFD performance comparison between our proposed HUGE and baselines with AUROC and AUPRC. The best and second-best results are in **bold** and underlined, respectively. OOM indicates out-of-memory on a 24GB GPU.

In detail, we first propagate and project the output embedding  $\mathbf{E}$  from the MLP encoder using a GNN layer:

$$\bar{\mathbf{E}} = \text{GNN}(\mathbf{E}, \mathbf{A}), \quad (9)$$

where  $\mathbf{A}$  is the adjacent matrix. After that, we align the NID (i.e. the distribution of rescaled cosine similarity between the embedding of  $v_i$  and its neighbors) of  $\mathbf{E}$  and  $\bar{\mathbf{E}}$  using an asymmetric alignment loss:

$$\begin{aligned} \mathcal{L}_{\text{align}} &= \frac{1}{|\mathcal{V}|} \sum_{v_i \in \mathcal{V}} \text{KLD}(s_i^{\text{NID}} | \text{sg}[\bar{s}_i^{\text{NID}}]) \\ &= \frac{1}{|\mathcal{V}|} \sum_{v_i \in \mathcal{V}} \frac{1}{|\mathcal{N}(v_i)|} \sum_{v_j \in \mathcal{N}(v_i)} \text{sg}[\bar{s}_{i,j}] (\log(\text{sg}[\bar{s}_{i,j}]) - \log(s_{i,j})), \end{aligned} \quad (10)$$

where  $\text{sg}[\cdot]$  denotes stop gradient operation,  $\text{KLD}(\cdot|\cdot)$  denotes the Kullback-Leibler Divergence, and  $s_{i,j}$ ,  $\bar{s}_{i,j}$  represent the rescaled cosine similarity (rescaled to  $[0, 1]$ ) between the embeddings of node  $i$  and  $j$  for MLP and GNN, respectively. The  $\mathcal{L}_{\text{align}}$  transfers the structural information captured by the GNN to the MLP by minimizing the KLD. This approach not only enhances the ability of the model to detect fraudulent entities, but also reduces the feature-smooth effect introduced by solely using GNN. Note that we employ an asymmetry loss with a stop gradient to prevent the GNN from degrading into an identical mapping.

Similar to the training of the MLP encoder, we also compute the ranking loss  $\mathcal{L}_{\text{rank}}$  for the GNN branch following Eq. (8), which results in our overall objective with a hyper-parameter  $\alpha$  to control the strength of alignment:

$$\mathcal{L} = \mathcal{L}_{\text{rank}} + \bar{\mathcal{L}}_{\text{rank}} + \alpha \mathcal{L}_{\text{align}}. \quad (11)$$

**Complexity Analysis.** To sum up, the heterophily estimation, training, and testing complexities of HUGE are  $\mathcal{O}(md)$ ,  $\mathcal{O}((nd+md_e+n|\mathcal{B}|)d_e)$  (per epoch), and  $\mathcal{O}((nd+m)d_e)$ , respectively, where  $d_e$  is the embedding dimension. Detailed complexity analysis can be found in Appendix D.

## Experiments

### Experimental Settings

**Datasets.** We conduct experiments on six public real-world GFD datasets, covering domains ranging from social networks to e-commerce, including Amazon (Dou et al. 2020), Facebook (Xu et al. 2022), Reddit, YelpChi (Kumar, Zhang,

and Leskovec 2019), AmazonFull (McAuley and Leskovec 2013), and YelpChiFull (Rayana and Akoglu 2015). Following (He et al. 2024; Qiao and Pang 2024), we convert all graphs to homogeneous undirected graphs in our experiments.

**Baselines and Evaluation Metrics.** We compare our method with 7 state-of-the-art (SOTA) unsupervised GAD methods: DOMINANT (Ding et al. 2019), CoLA (Liu et al. 2021b), ANEMONE (Jin et al. 2021), GRADATE (Duan et al. 2023a), ADA-GAD (He et al. 2024), GADAM (Chen et al. 2024a) and TAM (Qiao and Pang 2024). The area under the receiver operating characteristic curve (AUROC) and the area under the precision-recall curve (AUPRC) are used as the evaluation metrics. We report the average performance over five runs with different random seeds. The implementation details, hyper-parameters used in our experiments, and more results are presented in Appendix E and F.

### Experimental Results

**Performance Comparison.** The comparison results on six real-world GFD datasets are reported in Table 2. From the results, we make the following key observations: ① HUGE outperforms SOTA methods in most datasets in terms of both metrics, except for AUROC on Reddit and AUPRC on YelpChi. These results demonstrate the effectiveness of HUGE in diverse real-world GFD scenarios. Moreover, while existing methods show unstable performance across different datasets, HUGE consistently achieves competitive results. ② Among the baselines, the methods that consider heterophily (ADA-GAD, GADAM, TAM) usually outperform other baselines, demonstrating the benefits of mitigating heterophily for unsupervised GFD. However, their performances are constrained by their heterophily estimations. In contrast, our proposed heterophily measurement HALO shows expressive performance. ③ The scalability of several baselines is challenged by large-scale real-world GFD datasets with out-of-memory (OOM) issues, e.g., YelpChi-Full. In contrast, HUGE exhibits better scalability, highlighting its potential in handling real-world large-scale data from online services.

**Ablation Study.** To examine the contribution of key design in HUGE, we conduct ablation studies on heterophily measures and critical model components.

To evaluate the effectiveness of our label-free heterophily




Figure 3: Visualization of results. (a) Parameter sensitivity w.r.t.  $\alpha$  of the proposed HUGO. (b) Fraud score visualization on AmazonFull dataset. (c) AUROC and training time of HUGO and baselines on AmazonFull dataset.

Measures	Amazon	Facebook	AmazonFull	YelpChiFull
<b>HUGO (w/ HALO)</b>	<b>0.8516</b>	<b>0.9760</b>	<b>0.8892</b>	<b>0.5767</b>
w/ Euc. Dist.	0.7128	0.9251	0.8746	0.5763
w/ Cos. Dist.	0.7018	0.9668	0.8735	0.5762
w/ AHR	0.8171	0.9438	0.8622	0.5765
w/o Alignment	0.5614	0.9487	0.8746	0.5295
w/o GNN	0.5548	0.9401	0.7005	0.5533

Table 3: Ablation study of key designs in HUGO.

HALO (w/ HALO), we replace it with other heterophily measures, i.e. Euclidean distance (w/ Euc. Dist.), cosine distance (w/ Cos. Dist.), attribute heterophily rate (w/ AHR), as illustrated in the upper part of Table 3. We can observe that our HALO achieves superior performances compared to other heterophily measures, which verifies its effectiveness for unsupervised GFD. In addition, while some other measures demonstrate promising results on certain datasets, none of them consistently excels across all datasets, highlighting their limitations in guiding GFD models.

We compare HUGO with two variants that exclude the corresponding components, i.e., the asymmetric alignment loss (w/o Alignment) and GNN branch (w/o GNN), to verify the effectiveness of these key designs. The results in the lower part of Table 3 demonstrate that HUGO, with all components, achieves the best performance among all datasets, illustrating the effectiveness and robustness of each module. Meanwhile, the asymmetric alignment loss provides the major contribution, when removing it results in a significant performance drop. Moreover, while solely attaching the GNN model may lead to a performance drop (as illustrated in YelpChiFull) due to feature-smoothing effects, our alignment loss can overcome the issue by extracting the structural information to the MLP encoder.

**Sensitivity Analysis.** In order to evaluate the sensitivity of HUGO to the hyper-parameter  $\alpha$ , we adjust its value across  $\{0.0, 0.5, 1.0, 1.5, 2.0\}$ . The results are illustrated in Figure 3(a). Overall, HUGO is not sensitive to variation in  $\alpha$  on Facebook and Reddit, and slightly sensitive on Amazon and YelpChi. Moreover, HUGO achieves the best per-

formance at  $\alpha = 0.5$ , with a noticeable drop when  $\alpha = 0$  on Amazon and Facebook. These findings suggest HUGO performs best when neighbor information and ego information are balanced. Additionally, the small performance gap between  $\alpha = 0$  and  $\alpha = 1$  on Reddit and YelpChi shows that these datasets require less neighbor information for effective fraud detection.

**Visualization.** The distribution of learned fraud scores is visualized in Figure 3(b) by plotting a subset of scores. We observe that the fraudsters tend to have higher fraud scores, while benign nodes receive lower scores. The gap between the distributions of fraud and benign nodes highlights the effectiveness of HUGO. Moreover, the score distribution of fraudsters is more diverse than the score distribution of normal nodes, indicating the diversity of frauds in real-world unsupervised GFD scenarios. Nevertheless, HUGO still achieves favorable results.

**Efficiency Comparison.** To compare the efficiency of our method against baselines, we plot their training time against AUROC. As shown in Figure 3(c), HUGO obtains the best performance while maintaining superior efficiency. Although DOMINANT and GADAM can be trained faster, they fail to effectively detect fraudsters, performing worse than random guesses (AUROC=0.5). The high running efficiency indicates the potential of applying HUGO for real-time detection in practical scenarios.

## Conclusion

In this paper, we propose a novel unsupervised graph fraud detection (GFD) method termed HUGO, which effectively alleviates the negative effects of heterophily introduced by fraudsters. To address the challenge of estimating heterophily without labeled annotations, we design an effective label-free heterophily metric, HALO. Then, we use HALO to guide the training of our proposed model through a ranking loss. Additionally, we introduce an asymmetric alignment loss to extract useful structural information while mitigating the smoothing effect introduced by GNNs. Comprehensive experiments showcase the superior performance of HUGO over baseline methods across diverse datasets.

## Acknowledgments

This research was partly funded by Australian Research Council (ARC) under grants FT210100097 and DP240101547 and the CSIRO – National Science Foundation (US) AI Research Collaboration Program.

## References

- Burges, C.; Shaked, T.; Renshaw, E.; Lazier, A.; Deeds, M.; Hamilton, N.; and Hullender, G. 2005. Learning to rank using gradient descent. In *Proceedings of the 22nd international conference on Machine learning*, 89–96.
- Cai, C.; and Wang, Y. 2020. A note on over-smoothing for graph neural networks. *arXiv preprint arXiv:2006.13318*.
- Cai, J.; and Fan, J. 2022. Perturbation learning based anomaly detection. In *Advances in Neural Information Processing Systems*, 14317–14330.
- Cai, J.; Zhang, Y.; Fan, J.; and Ng, S.-K. 2024. LG-FGAD: An Effective Federated Graph Anomaly Detection Framework. In *Proceedings of the International Joint Conference on Artificial Intelligence*.
- Chen, J.; Zhu, G.; Yuan, C.; and Huang, Y. 2024a. Boosting Graph Anomaly Detection with Adaptive Message Passing. In *The Twelfth International Conference on Learning Representations*.
- Chen, N.; Liu, Z.; Hooi, B.; He, B.; Fathony, R.; Hu, J.; and Chen, J. 2024b. Consistency Training with Learnable Data Augmentation for Graph Anomaly Detection with Limited Supervision. In *The Twelfth International Conference on Learning Representations*.
- Deng, L.; Wu, C.; Lian, D.; Wu, Y.; and Chen, E. 2022. Markov-driven graph convolutional networks for social spammer detection. *IEEE Transactions on Knowledge and Data Engineering*, 35(12): 12310–12322.
- Ding, K.; Li, J.; Bhanushali, R.; and Liu, H. 2019. Deep anomaly detection on attributed networks. In *Proceedings of the 2019 SIAM international conference on data mining*, 594–602. SIAM.
- 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.
- Duan, J.; Wang, S.; Zhang, P.; Zhu, E.; Hu, J.; Jin, H.; Liu, Y.; and Dong, Z. 2023a. Graph anomaly detection via multi-scale contrastive learning networks with augmented view. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, 7459–7467.
- Duan, J.; Xiao, B.; Wang, S.; Zhou, H.; and Liu, X. 2023b. Arise: Graph anomaly detection on attributed networks via substructure awareness. *IEEE transactions on neural networks and learning systems*.
- 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.
- Gu, J.; and Zou, D. 2024. Three Revisits to Node-Level Graph Anomaly Detection: Outliers, Message Passing and Hyperbolic Neural Networks. In *Learning on Graphs Conference*, 14–1. PMLR.
- He, J.; Xu, Q.; Jiang, Y.; Wang, Z.; and Huang, Q. 2024. ADA-GAD: Anomaly-Denoised Autoencoders for Graph Anomaly Detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 8481–8489.
- 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.
- Kumar, S.; Zhang, X.; and Leskovec, J. 2019. Predicting dynamic embedding trajectory in temporal interaction networks. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, 1269–1278.
- Li, S.; Liu, Y.; Chen, Q.; Webb, G. I.; and Pan, S. 2024. Noise-Resilient Unsupervised Graph Representation Learning via Multi-Hop Feature Quality Estimation. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*.
- Li, Z.; Liu, Y.; Zhang, Z.; Pan, S.; Gao, J.; and Bu, J. 2022. Cyclic label propagation for graph semi-supervised learning. *World Wide Web*, 25(2): 703–721.
- Liu, Y.; Ajanthan, T.; Husain, H.; and Nguyen, V. 2024a. Self-supervision improves diffusion models for tabular data imputation. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*.
- 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 *Proceedings of the web conference 2021*, 3168–3177.
- Liu, Y.; Ding, K.; Lu, Q.; Li, F.; Zhang, L. Y.; and Pan, S. 2023. Towards self-interpretable graph-level anomaly detection. *Advances in Neural Information Processing Systems*, 36.
- Liu, Y.; Li, S.; Zheng, Y.; Chen, Q.; Zhang, C.; and Pan, S. 2024b. ARC: A Generalist Graph Anomaly Detector with In-Context Learning. In *Advances in Neural Information Processing Systems*.
- Liu, Y.; Li, Z.; Pan, S.; Gong, C.; Zhou, C.; and Karypis, G. 2021b. Anomaly detection on attributed networks via contrastive self-supervised learning. *IEEE transactions on neural networks and learning systems*, 33(6): 2378–2392.
- 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.

- Ma, X.; Wu, J.; Xue, S.; Yang, J.; Zhou, C.; Sheng, Q. Z.; Xiong, H.; and Akoglu, L. 2021. A comprehensive survey on graph anomaly detection with deep learning. *IEEE Transactions on Knowledge and Data Engineering*, 35(12): 12012–12038.
- 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.
- Motie, S.; and Raahemi, B. 2023. Financial fraud detection using graph neural networks: A systematic review. *Expert Systems With Applications*, 122156.
- Pan, J.; Liu, Y.; Zheng, Y.; and Pan, S. 2023. PREM: A Simple Yet Effective Approach for Node-Level Graph Anomaly Detection. In *2023 IEEE International Conference on Data Mining (ICDM)*, 1253–1258. IEEE.
- Platonov, O.; Kuznedelev, D.; Babenko, A.; and Prokhorenkova, L. 2024. Characterizing graph datasets for node classification: Homophily-heterophily dichotomy and beyond. *Advances in Neural Information Processing Systems*, 36.
- Qiao, H.; and Pang, G. 2024. Truncated affinity maximization: One-class homophily modeling for graph anomaly detection. In *Advances in Neural Information Processing Systems*, volume 36.
- 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.
- 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.
- Wang, L.; Zheng, Y.; Jin, D.; Li, F.; Qiao, Y.; and Pan, S. 2024a. Contrastive graph similarity networks. *ACM Transactions on the Web*, 18(2): 1–20.
- Wang, Y.; Liu, Y.; Shen, X.; Li, C.; Ding, K.; Miao, R.; Wang, Y.; Pan, S.; and Wang, X. 2024b. Unifying Unsupervised Graph-Level Anomaly Detection and Out-of-Distribution Detection: A Benchmark. *arXiv preprint arXiv:2406.15523*.
- 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.
- Xu, Z.; Huang, X.; Zhao, Y.; Dong, Y.; and Li, J. 2022. Contrastive attributed network anomaly detection with data augmentation. In *Pacific-Asia conference on knowledge discovery and data mining*, 444–457. Springer.
- Yang, L.; Li, M.; Liu, L.; Wang, C.; Cao, X.; Guo, Y.; et al. 2021. Diverse message passing for attribute with heterophily. *Advances in Neural Information Processing Systems*, 34: 4751–4763.
- Yu, S.; Ren, J.; Li, S.; Naseriparsa, M.; and Xia, F. 2022. Graph learning for fake review detection. *Frontiers in Artificial Intelligence*, 5: 922589.
- Zhang, Y.; Sun, Y.; Cai, J.; and Fan, J. 2024. Deep Orthogonal Hypersphere Compression for Anomaly Detection. In *Proceedings of the Twelfth International Conference on Learning Representations*.
- Zhao, L.; and Akoglu, L. 2020. On Using Classification Datasets to Evaluate Graph Outlier Detection: Peculiar Observations and New Insights. *CoRR*, abs/2012.12931.
- Zheng, X.; Song, D.; Wen, Q.; Du, B.; and Pan, S. 2024a. Online GNN Evaluation Under Test-time Graph Distribution Shifts. In *The Twelfth International Conference on Learning Representations*.
- Zheng, X.; Wang, Y.; Liu, Y.; Li, M.; Zhang, M.; Jin, D.; Yu, P. S.; and Pan, S. 2022a. Graph neural networks for graphs with heterophily: A survey. *arXiv preprint arXiv:2202.07082*.
- Zheng, X.; Zhang, M.; Chen, C.; Molaei, S.; Zhou, C.; and Pan, S. 2024b. Gnnevaluator: Evaluating gnn performance on unseen graphs without labels. *Advances in Neural Information Processing Systems*, 36.
- Zheng, Y.; Jin, M.; Liu, Y.; Chi, L.; Phan, K. T.; and Chen, Y.-P. P. 2021a. Generative and contrastive self-supervised learning for graph anomaly detection. *IEEE Transactions on Knowledge and Data Engineering*, 35(12): 12220–12233.
- Zheng, Y.; Lee, V. C.; Wu, Z.; and Pan, S. 2021b. Heterogeneous graph attention network for small and medium-sized enterprises bankruptcy prediction. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, 140–151. Springer.
- Zheng, Y.; Pan, S.; Lee, V.; Zheng, Y.; and Yu, P. S. 2022b. Rethinking and scaling up graph contrastive learning: An extremely efficient approach with group discrimination. *Advances in Neural Information Processing Systems*, 35: 10809–10820.
- Zheng, Y.; Zhang, H.; Lee, V.; Zheng, Y.; Wang, X.; and Pan, S. 2023. Finding the missing-half: Graph complementary learning for homophily-prone and heterophily-prone graphs. In *International Conference on Machine Learning*, 42492–42505. PMLR.
- Zheng, Y.; Zheng, Y.; Zhou, X.; Gong, C.; Lee, V. C.; and Pan, S. 2022c. Unifying graph contrastive learning with flexible contextual scopes. In *2022 IEEE International Conference on Data Mining (ICDM)*, 793–802. IEEE.
- Zhong, Q.; Liu, Y.; Ao, X.; Hu, B.; Feng, J.; Tang, J.; and He, Q. 2020. Financial defaulter detection on online credit payment via multi-view attributed heterogeneous information network. In *Proceedings of the web conference 2020*, 785–795.
- Zhu, J.; Jin, J.; Loveland, D.; Schaub, M. T.; and Koutra, D. 2022. How does heterophily impact the robustness of graph neural networks? theoretical connections and practical implications. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2637–2647.
- Zhuo, W.; Liu, Z.; Hooi, B.; He, B.; Tan, G.; Fathony, R.; and Chen, J. 2024. Partitioning message passing for graph fraud detection. In *The Twelfth International Conference on Learning Representations*.