Dynamic Neighborhood Modeling via Node-Subgraph Contrastive Learning for Graph-Based Fraud Detection

Zhizhi Yu[1], Chundong Liang[1,2], Xinglong Chang[1,3], Dongxiao He[1], Di Jin[1]*, Jianguo Wei[1]*

[1]College of Intelligence and Computing, Tianjin University, Tianjin, China
[2]North China University of Science and Technology, Tangshan, Hebei, China
[3]Qijia Youdao Network Technology (Beijing) Co., Ltd., Beijing, China
{yuzhizhi, liangchundong, changxinglong, hedongxiao, jindi, jianguo}@tju.edu.cn

Abstract

Fraud detection that aims to discern frauds from the majority of benigns has become an increasingly prominent research field. Recently, Graph Neural Networks (GNNs) have been widely applied in graph-based fraud detection due to their outstanding data analysis and mining capabilities. However, owing to the inherent homophily-heterophily mixture and class imbalance of fraud graphs, most GNNs with homophily assumption inevitably suffer from local abnormal signal loss during information propagation, posing significant challenges in situations where frauds are rare and valuable. To address the aforementioned issues, we present a novel dynamic neighborhood modeling via node-subgraph contrastive learning for graph-based fraud detection, dubbed DCL-GFD. Specifically, we first design a node abnormality estimation module from the perspective of feature, which analyses the likelihood of a node belonging to fraud or benign by comparing the feature similarity between the target node and its corresponding subgraph. We then present a dynamic neighborhood modeling mechanism guided by the abnormal probability of a node to adaptively group and aggregate neighborhood information. By this means, the target node can effectively aggregate the neighbor information from the perspective of fraud or benign, thereby preserving as much fraud characteristics that occupy minority population as possible. Extensive experiments across four real-world fraud detection datasets demonstrate the superiority and effectiveness of our proposed DCL-GFD over state-of-the-art baselines.

Introduction

With the rapid development of the Internet economy and beyond networks, various types of fraud activities have increased significantly, such as deceptive comments, illicit bitcoin transactions and malicious websites, making the detection of such activities a crucial research area in both academia and industry. In recent years, graph-based fraud detection approaches (Zhang et al. 2021; Wu et al. 2023; Dong et al. 2024) have become a promising development by modeling the entities as nodes and the corresponding interactions between entities as edges. In this way, fraud detection can be considered as a binary node classification task (Tang et al. 2022; Gao et al. 2024), whereby the objective is to distinguish whether a target node is fraud or benign. For example, by closely analyzing financial transaction graphs, we can detect potential suspicious patterns and thus identify frauds, e.g., nodes that exhibit frequent large transactions with multiple different accounts.

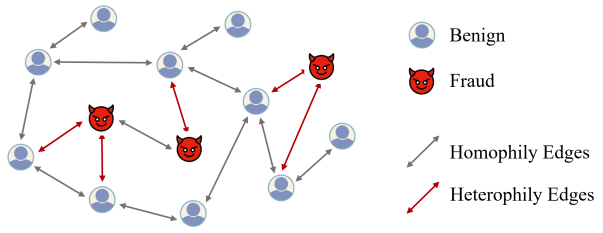

Figure 1: A toy example of fraud graph with homophily-heterophily mixture and class imbalance.

Along with the lastest advances of graph neural networks (GNNs) in graph representation learning, a great quantity of GNNs with diverse architectures have paved a new way for discriminating frauds (Wang et al. 2022; Yu et al. 2023). Typical GNNs are usually based on the homophily assumption (Wang et al. 2020; Wu et al. 2021), where node embeddings are obtained through feature propagates along graph topology. However, fraud detection scenarios (see an example in Figure 1) are essentially a mixture of homophily and heterophily (Gao et al. 2023a; Zhuo et al. 2024). That is, fraud nodes are typically submerged themselves in benign communities through structural or feature camouflage to mitigate suspiciousness, exhibiting heterophily, while the context surrounding benign nodes exhibits homophily. This makes existing GNNs inevitably suffer from local abnormal signal loss during information propagation, thereby compromising their ability to precisely distinguish frauds. More importantly, class imbalance is commonly present in the domain of fraud detection (Wang et al. 2023), leading GNNs to primarily capture the patterns and characteristics of benign nodes. That is to say, compared to a large number of benign nodes, it is of great importance to preserve the patterns and characteristics of fraud nodes that are rare and informative.

Several algorithms and models have been designed to diminish the impact of homophily-heterophily mixture and class imbalance during information propagation. For example, from the perspective of spectral domain, it has been attempted to associate graph anomalies with high frequency

*Corresponding authors.

spectral distributions, and accordingly construct GNN filters in different frequency bands to facilitate fraud detection (Chai et al. 2022; Tang et al. 2022). However, these models struggle with complex spectral decomposition, facing scalability issues on large-scale graphs due to the explosive memory requirements. Another line of attempts is spatial-centric approach, which usually centers around formulating models by closely analyzing the neighborhood of nodes to be classified, e.g., reweighting or resampling neighbors for the target node (Shi et al. 2022), so as to alleviate the impact of heterophily and imbalanced class distributions. In particular, PMP (Zhuo et al. 2024) designs a new neighbor aggregation strategy that restricts parameter sharing to nodes within the same class, including fraud, benign and unlabeled. CONSIS-GAD (Chen et al. 2024) proposes a learnable data augmentation method that is anchored in the principles of consistency training to discriminate fraud or benign nodes. However, these methods suffer from an inability to essentially exploit the abnormal tendency of neighbor nodes to guide information propagation, as they mainly distinguish neighbors via straightforward ways such as learning different weight matrices for homophily and heterophily edges separately.

In light of the aforementioned issues, we present a new **D**ynamic neighborhood modeling via node-subgraph **C**ontrastive **L**earning for **G**raph-based **F**raud **D**etection, termed DCL-GFD. Specifically, we first design a node-subgraph contrastive principle that integrates partially observed node labels from a feature perspective to estimate the abnormal probability of nodes. That is, if the feature similarity between a node and its corresponding subgraph is relatively high, it is more likely to be benign, whereas it tends to be fraud. We then introduce a dynamic neighborhood modeling mechanism continuously guided by the abnormal tendency of neighbor nodes, which aggregates neighbor information from fraud-like and benign-like perspectives, respectively. In this way, the model can essentially capture the uncertain and mixed nature of neighbors associated with the target node, thereby effectively relieving the inherent homophily-heterophily mixture and class imbalance within fraud detection scenarios.

The main contributions are summarized as follows:

- We gain a deep insight into the fraud detection via essentially considering the abnormal tendency of nodes themselves to preserve the informative patterns and characteristics of fraud nodes.
- We propose a new GNN-based fraud detection approach DCL-GFD, which introduces dynamic neighborhood modeling from the perspective of fraud or benign under the guidance of feature comparison of nodes and their corresponding subgraphs.
- Extensive experiments are conducted on four public fraud detection datasets, demonstrating the superior effectiveness and robustness of the proposed DCL-GFD over state-of-the-art baselines.

Preliminaries

We first present the notations and problem definition, and then introduce the backbone graph neural networks.

Notations and Problem Definition

Attributed Graph. Consider an attributed graph $G = (\mathcal{V}, \mathcal{E}, X, Y)$, with a set of N nodes $\mathcal{V} = \{v_1, \ldots, v_N\}$, a set of edges $\mathcal{E} = \{e_{ij}\} \subseteq \mathcal{V} \times \mathcal{V}$ describing the relations among nodes, a node feature matrix X wherein the i-th row x_i is the feature vector of v_i, and a set of corresponding labels Y of the nodes. The topological structure of G is represented by an adjacency matrix $A = [a_{ij}] \in \{0, 1\}^{N \times N}$, where $a_{ij} = 1$ if node v_i connects to node v_j, otherwise $a_{ij} = 0$.

Subgraph Sampling. For a target node v_i, the corresponding subgraph is typically defined as the adjacent substructure near the node (Duan et al. 2023b). Recent researches have shown that the feature similarity between the node and its subgraph can be effectively utilized to estimate its abnormality (Duan et al. 2023a). To this end, we adopt random walk with restart (RWR) (Qiu et al. 2020) to sample the k-step subgraph around the target node with restart probability r due to its usability. It is worth noting that a lower feature similarity indicates a higher probability that the target node tends to be a fraud.

Graph-Based Fraud Detection. Given an attributed graph G, the graph-based fraud detection (GFD) can be conceptualized as a node-level imbalanced binary classification problem, wherein each node should be predicted as either benign (the majority class) associated with a label of 0, or fraud (the minority class) related to a label of 1. Mathematically, let $\mathcal{V}_{\mathcal{L}} \ll N$ be a set of labeled nodes, the objective of GFD is to learn a mapping function $\mathcal{F}$ to calculate the labels of the remaining nodes $\mathcal{V} \backslash \mathcal{V}_{\mathcal{L}}$.

Graph Neural Networks

Classical graph neural networks (GNNs) and their variants (Kipf and Welling 2017) typically adopt a message passing principle, where the most essential part is the feature propagates along graph topology. Let $h_i^{(l)}$ be the hidden embedding of node v_i at l-th layer, the general framework of a GNN model can be calculated as:

$$h_i^{(l+1)} = h_i^{(l)} \oplus \mathrm{AGG}^{(l)}(\{h_j^{(l)} : v_j \in \mathcal{N}_i\}), \qquad (1)$$

where $h_i^{(0)}$ is the initial feature vector of node v_i, $\mathcal{N}_i$ denotes the neighbors of node v_i, AGG is the aggregation function that aggregates information from neighbors, and $\oplus$ represents the combination of the target node information and its neighbor information. GNNs perform well in learning node embeddings and applying them to different analytical tasks (Yun et al. 2021), but graph-based fraud detection suffers from serious class imbalance and heterophily problems, posing unique challenges to existing GNNs.

Methodology

We begin with a brief overview of the proposed method, followed by a detailed description of each component.

Overview

To effectively characterize the informativeness of neighbors during the propagation process of GNNs and perfectly align it with the ultimate goal of graph-based fraud detection, we propose a dynamic neighborhood modeling through

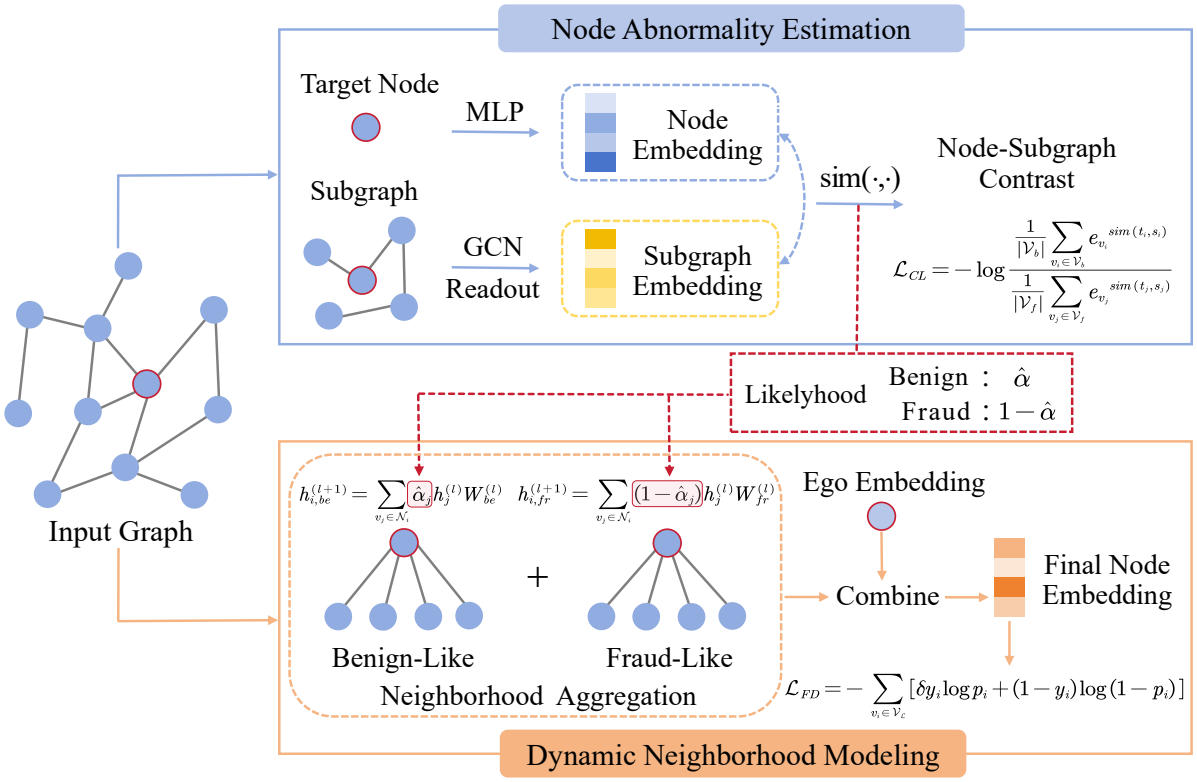

Figure 2: A sketch of DCL-GFD, which introduces a dynamic neighborhood modeling that continuously refines via node-subgraph contrastive learning for graph-based fraud detection.

node-subgraph contrastive learning, namely DCL-GFD. The whole architecture of the proposed approach is displayed in Figure 2, which comprises two principal components: node abnormality estimation as well as dynamic neighborhood modeling. Specifically, we first estimate node abnormality from the feature comparison of nodes and their corresponding subgraphs, wherein the common sense is that benigns should have higher feature similarity with their subgraphs, whereas frauds have lower feature similarity. We then design a dynamic grouping neighborhood modeling mechanism continuously guided by the abnormal probability of neighbor nodes, which aggregates neighbor information from fraud and benign perspectives respectively. In this way, the target node can adaptively adjust the influence of neighbors, thereby alleviating the problem of homophily-heterophily mixture, which is of significant importance for frauds that occupy minority population.

Node Abnormality Estimation

Inspired by the effectiveness of contrastive learning paradigm in local inconsistency mining (Jin et al. 2021; Duan et al. 2023b), we first estimate the abnormal probability of nodes through node-subgraph contrast from the perspective of feature. That is, if the feature similarity between a node and its corresponding subgraph is relatively high, it tends to be a benign; whereas if the feature similarity is relatively low, it is more likely to be a fraud.

Following (Duan et al. 2023a), for a target node v_i, we sample its corresponding subgraph G_i with adjacency matrix A_i and attribute matrix X_i using random walk with restart. Then, we adopt graph convolutional network (GCN) (Kipf and Welling 2017) to learn node embeddings within subgraphs, where the feature of the target node are anonymized by setting their value to 0 to prevent the influence of the node itself. Formally, the hidden layer embedding of the subgraph can be expressed as:

$$H_i^{(l+1)} = \sigma(\tilde{D}_i^{-\frac{1}{2}}\tilde{A}_i\tilde{D}_i^{-\frac{1}{2}}H_i^{(l)}W^{(l)}), \tag{2}$$

where $H_i^{(0)} = X_i$, $\tilde{D}_i^{-\frac{1}{2}}\tilde{A}_i\tilde{D}_i^{-\frac{1}{2}}$ stands for the normalization of the adjacency matrix with $\tilde{D}_i = \text{diag}\ (\tilde{d}_1, ..., \tilde{d}_N)$ and $\tilde{d}_i = \sum_j \tilde{a}_{ij}$, $W^{(l)}$ denotes the trainable weight matrix at $(l+1)$-th layer, and σ denotes the non-linear activation function such as ReLU.

After obtaining the embedding of the subgraph S_i, we convert it to the same shape as target node v_i by introducing a Readout function. Here we utilize the average function to achieve Readout, and the final embedding of the subgraph s_i can then be defined as:

$$s_i = \text{Readout}(S_i) = \frac{1}{\tilde{n}_i}\sum_{j=1}^{\tilde{n}_i}(S_i)_j, \tag{3}$$

where $\tilde{n}_i$ represents the number of nodes within the corresponding subgraph of target node v_i, and $(S_i)_j$ denotes the j-th row of S_i.

In the meantime, to ensure that the target node embedding lies in the same embedding space with the subgraph embedding, we introduce a multi-layer feed forward network (FFN) to calculate the hidden embedding of the target node, calculated as:

$$h_i^{(l+1)} = \sigma(h_i^{(l)} W^{(l)}), \tag{4}$$

where $W^{(l)}$ and σ denote the trainable weight matrix and activation function, respectively. In this way, we can obtain the final target node embedding t_i.

By calculating the similarity between the target node embedding t_i and its corresponding subgraph embedding s_i, we can obtain the abnormal probability α of the target node. Here we adopt cosine similarity that calculates the similarity utilizing the cosine value of angle among embeddings to estimate node abnormality as:

$$\alpha_i = \mathrm{sim}(t_i, s_i) = \frac{t_i \cdot s_i}{|t_i||s_i|}. \tag{5}$$

A larger α_i means that the target node is more similarly to a benign, and vice versa.

Finally, considering that the advantages of label information in enhancing the reliability of the abnormal probability of a node (Xu et al. 2024), inspired by the InfoNCE loss (You et al. 2020) in contrastive learning, we further design a node-subgraph contrastive loss function that integrates partially observed node labels. Specifically, the embeddings of a benign and its corresponding subgraph should be as similar as possible, whereas the two embeddings of a fraud should have a large deviation, represented as:

$$\mathcal{L}_{CL} = -\log \frac{\frac{1}{|\mathcal{V}_b|}\sum_{v_i \in \mathcal{V}_b} e^{\alpha_i}_{v_i}}{\frac{1}{|\mathcal{V}_f|}\sum_{v_j \in \mathcal{V}_f} e^{\alpha_j}_{v_j}}, \tag{6}$$

where $\mathcal{V}_b$ and $\mathcal{V}_f$ denote the set of benign nodes and fraud nodes in the training set, respectively. It is worth noting that due to the randomness of subgraph sampling strategy, it is difficult to reflect all feature information of neighbors in one sampling. To this end, we perform multiple subgraph sampling during model training.

Dynamic Neighborhood Modeling

Owing to the prevalence of class imbalance and heterophily issues in fraud detection scenarios, traditional graph neural networks with homophily assumption are prone to result in local abnormal signal loss during information propagation (Wang et al. 2023; Zhuo et al. 2024). To effectively distill valuable information related to target node detection from neighbors, we design a dynamic neighborhood modeling mechanism that aggregates neighbor information based on the guidance of the probability that neighbors belonging to benigns or frauds, respectively. This essentially models the uncertain and mixed nature of neighbors associated with the target node by modulating their class tendency.

Specifically, for a target node v_i, we first aggregate its neighbor information from the perspective of the probability that neighbors are more inclined to be benigns. Thus, the benign-like neighbor embedding is defined as:

$$h_{i,be}^{(l+1)} = \sum_{v_j \in \mathcal{N}_i} \hat{\alpha}_j h_j^{(l)} W_{be}^{(l)}, \tag{7}$$

where $\hat{\alpha}_j$ describes the probability that neighbor node v_j belongs to a benign node, which is normalized by α_j, i.e., $\hat{\alpha}_j = (1+\alpha_j)/2$. $l \in \{0, \cdots, L-1\}$ denotes the layer index, and $W_{be}^{(l)}$ represents the $(l+1)$-th layer weight matrix associated with benign, which is shared across all target nodes when performing information propagation.

Analogously, we obtain the fraud-like neighbor embedding utilizing the probability that neighbors are more prone to be frauds as:

$$h_{i,fr}^{(l+1)} = \sum_{v_j \in \mathcal{N}_i} (1-\hat{\alpha}_j) h_j^{(l)} W_{fr}^{(l)}, \tag{8}$$

where $W_{fr}^{(l)}$ represents the $(l+1)$-th layer weight matrix related to fraud.

Then, by integrating the aforementioned benign-like and fraud-like neighbor embeddings, the hidden layer neighbor embedding of node v_i can be expressed as:

$$h_{i,nei}^{(l+1)} = h_{i,be}^{(l+1)} + h_{i,fr}^{(l+1)}. \tag{9}$$

In this way, the target node can dynamically model neighbor knowledge based on the learned abnormal probability of neighbor nodes during information propagation, so as to relieve class imbalance and heterophily issues, and thus increase the discriminative ability of nodes.

Considering that ego and neighbor embeddings are likely to be dissimilar of frauds due to heterophily (Zhu et al. 2020), for target node v_i, we further concatenate its ego and neighbor embeddings to enable node to learn better embeddings from neighbors while also retaining its initial features to a certain extent, which is calculated as:

$$h_i^{(l+1)} = h_i^{(l)} \,\|\, h_{i,nei}^{(l+1)}, \tag{10}$$

where $\|$ denotes the concatenation function.

Finally, we feed the final layer's embedding $h_i^{(L)}$ into a multi-layer perceptron (MLP) followed by a sigmoid function to predict the probability that node v_i tends to be a fraud, elaborated as:

$$p_i = \mathrm{sigmoid}\left(\mathrm{MLP}(h_i^{(L)})\right). \tag{11}$$

Following the semi-supervised setting, the loss function of fraud detection can be defined by adopting weighted cross-entropy (Zhong, Wang, and Miao 2019):

$$\mathcal{L}_{FD} = -\sum_{v_i \in \mathcal{V}_\mathcal{L}} \left[\delta y_i \log p_i + (1-y_i)\log(1-p_i)\right], \tag{12}$$

where δ is the ratio of fraud labels ($y_i = 1$) to benign labels ($y_i = 0$) in the training set.

Model Optimization

For the sake of fully utilizing the learnable target node abnormal probability to dynamically guide the propagation of neighbor information, we unify and jointly optimize the node abnormality estimation module and dynamic neighborhood modeling module in an end-to-end manner. Thus, the final loss function can be defined as:

$$\mathcal{L} = \mathcal{L}_{FD} + \lambda \mathcal{L}_{CL}, \tag{13}$$

where λ serves as a hyperparameter to balance the influence of two parts of loss on model performance.

Experiments

We first introduce the experimental setup, and then evaluate the new approach DCL-GFD with state-of-the-art baselines to demonstrate its effectiveness. After that, we analyze the impact of different training ratios on model performance, as well as conduct a qualitative analysis of node abnormality estimation and parameter sensitivity.

Experimental Setup

Datasets. We conduct experiments on four real-world datasets targeted at fraud detection scenarios, that is, Weibo (Zhao et al. 2020), Amazon (McAuley and Leskovec 2013), YelpChi (Rayana and Akoglu 2015) and T-Finance (Tang et al. 2022). Specifically, Weibo seeks to identify anomalous accounts on social media platforms, where suspicious activities are defined as two posts made within a specific timeframes. Amazon aims to detect users who are compensated for producing deceptive product reviews within the musical instrument category. YelpChi focuses on identifying reviews that unfairly promote or demote specific products or businesses on Yelp.com. T-Finance aims to detect anomalous accounts within transaction graphs. The statistics of these datasets are shown in Table 1.

Datasets	#Nodes	#Edges	#Features	Fraud (%)
Weibo	8,405	407,963	400	10.30
Amazon	11,944	4,398,392	25	6.87
YelpChi	45,954	3,846,979	32	14.53
T-Finance	39,357	21,222,543	10	4.58

Table 1: Statistics of datasets.

Baselines. We evaluate the performance of DCL-GFD by comparing it with three categories of state-of-art algorithms. They include: 1) Classical GNNs, including GCN (Kipf and Welling 2017), GAT (Velickovic et al. 2018) and GraphSage (Hamilton, Ying, and Leskovec 2017), which work under the assumption of homophily; 2) Spectral GNNs for Fraud Detection, including AMNet (Chai et al. 2022), BWGNN (Tang et al. 2022) and GHRN (Gao et al. 2023a), which focus on recognizing the graph signals generated by frauds or benigns; 3) Spatial GNNs for Fraud Detection, including CARE-GNN (Dou et al. 2020), PC-GNN (Liu et al. 2021), GDN (Gao et al. 2023b), PMP (Zhuo et al. 2024), and CONSISGAD (Chen et al. 2024), which concentrate on designing distinct information propagation strategies to identify homophily or heterophily edges.

Implementation Details. We implement GCN, GAT and GraphSage in DGL (Wang et al. 2019), and for other baseline methods, we adopt the source code provided in the original paper. For our proposed DCL-GFD, the parameters are optimized by Adam, while the learning rate is set as 0.01, and the weight decay rate is 5e-3. We set the propagation layer of node abnormality estimation and dynamic neighborhood modeling to 1, the hidden embedding size to 64, and the batch size to 1024 on all four datasets. In addition, the step size and restart probability of random walk with restart used in subgraph sampling are set to 8 and 0.5, respectively. For a fair comparison of all methods, we set training, validation, and testing in a 4:2:4 ratio, and randomly run them 5 times to report the average results.

Evaluation Metrics. Considering that graph-based fraud detection poses a class-imbalanced classification problem, we employ two widely adopted metrics, AUC and Macro-F1, to measure the model performance.

Performance Comparison

We compare the performance of our proposed DCL-GFD with representative baselines in a fraud detection task, and report the mean Accuracy and Macro-F1 along with the standard deviation of 5 independent runs.

As presented in Table 2, DCL-GFD outperforms all baseline methods on 3 out of 4 datasets. Specifically, in terms of AUC, DCL-GFD achieves up to 2.03%, 0.07%, and 0.28% better accurate than the best baseline method on Weibo, Amazon, and YelpChi, respectively. In terms of Macro-F1, DCL-GFD is 2.62%, 0.73%, and 0.95% more accurate than the best baseline method on these three datasets. These results not only demonstrate the superiority of designing a node-subgraph contrast to estimate node abnormality from the view of feature, but also prove the significance of our new dynamic neighborhood modeling mechanism that simulates benign-like and fraud-like nature of neighbors related to the target node, respectively. It is worth noting that our DCL-GFD performs the second best on T-Finance dataset, which may be due to its limited initial feature information, affecting node abnormality estimation to a certain extent. In addition, compared with GraphSage, a representative graph neural network with homophily assumption, DCL-GFD attains improvements of 13.06%, 2.77%, 15.76% and 6.83% in AUC, and 6.33%, 1.99%, 17.27%, and 8.29% in Macro-F1 on Weibo, Amazon, YelpChi, and T-Finance, respectively. The reason behind this can be attributed to that our method effectively preserves as much fraud characteristics that occupy minority population as possible during the process of information propagation.

Quantitative Analysis of Training Ratio

Considering that annotating samples in fraud detection scenarios is a costly endeavor, we further investigate the performance of our proposed DCL-GFD under different training ratios. Specifically, we compare DCL-GFD with two SOTA baselines (i.e., PMP and CONSISGAD) on Weibo and Amazon datasets with training ratios ranging from 10% to 40%.

The analysis results are illustrated in Figure 3. As shown, we observe that DCL-GFD performs consistently the best compared with PMP and CONSISGAD across different training ratios, showcasing its efficacy under limited supervision. This is due to the fact that our DCL-GFD adopts dynamic neighborhood modeling to distill valuable information related to target node detection from neighbors, making the learned representations more discriminative. In addition, with a decrease in training data, the performance gap between our DCL-GFD and baselines gradually increase, especially on the Weibo dataset with 10% training ratio, This further demonstrates the robustness and effectiveness of our method under low training ratios.

Methods	Weibo		Amazon		YelpChi		T-Finance	
	AUC	Macro-F1	AUC	Macro-F1	AUC	Macro-F1	AUC	Macro-F1
GCN	96.16±0.39	92.71±0.35	80.41±0.82	63.62±1.57	54.68±1.36	52.23±0.98	80.27±2.18	68.87±3.03
GAT	94.76±1.85	88.84±1.91	84.28±4.18	63.58±3.12	52.97±1.70	51.42±1.47	82.66±0.84	73.85±6.28
GraphSage	85.93±1.66	87.74±0.26	95.05±1.42	91.17±0.41	77.15±1.23	64.43±0.81	89.75±1.08	82.48±3.41
AMNet	94.59±0.45	92.76±0.34	97.17±0.64	92.45±0.29	84.16±0.20	69.35±0.52	94.49±0.61	89.57±0.61
BWGNN	97.22±0.42	92.87±0.58	96.72±1.58	91.34±0.92	90.60±0.64	77.09±0.62	94.64±0.85	87.74±1.67
GHRN	96.16±0.45	90.90±0.58	96.51±0.63	91.92±0.43	90.64±0.34	77.42±0.61	95.84±0.37	88.41±1.22
CARE-GNN	89.34±3.54	78.83±2.34	93.75±0.51	86.89±1.10	77.53±0.04	61.77±1.65	91.31±0.19	71.88±1.51
PC-GNN	96.25±0.68	88.61±0.57	96.51±0.26	88.04±0.84	81.61±0.40	64.98±4.27	91.54±0.46	63.94±4.00
GDN	68.83±0.18	59.44±0.12	95.25±0.13	90.32±0.18	90.71±0.13	73.12±0.58	94.31±0.28	81.91±6.14
PMP	95.89±0.40	91.45±0.27	97.53±0.15	91.69±0.68	92.63±0.12	80.75±0.38	96.37±0.19	90.61±0.71
CONSISGAD	96.96±0.25	90.35±0.47	97.75±0.11	92.43±0.39	91.13±0.41	77.08±0.51	**96.87±0.21**	**91.20±0.39**
DCL-GFD	**98.99±0.05**	**94.07±0.33**	**97.82±0.15**	**93.16±0.23**	**92.91±0.11**	**81.70±0.10**	96.58±0.06	90.77±0.12

Table 2: Experiment results with mean value and standard deviation in terms of AUC (%) and Macro-F1 (%). Bold and underline are adopted to display the best and the second best results.

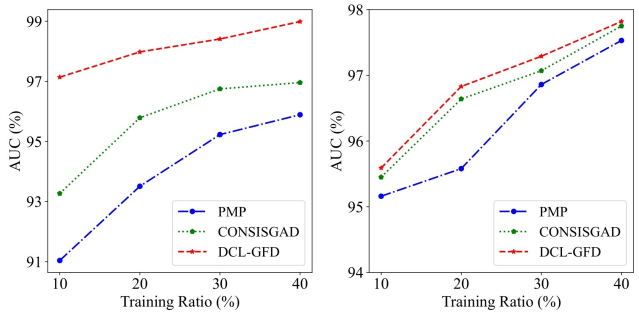

(a) Weibo (b) Amazon

Figure 3: Comparisons of our DCL-GFD with PMP and CONSISGAD under different training ratios in AUC (%).

Qualitative Analysis of Abnormality Estimation

We further assess and demonstrate the effectiveness of our proposed node abnormality estimation module using Amazon and YelpChi datasets as examples. Specifically, we randomly select 200 fraud nodes and 200 benign nodes for each dataset, and visualize the feature similarity between nodes and their corresponding subgraph embeddings. It is worth noting that in order to clarify the abnormal probability of a node more clearly, we reorganize the ID number of fraud nodes and benign nodes from 1 to 400, and the higher the feature similarity, the more likely the node is a benign, and vice versa, the node tends to be a fraud.

As shown in Figure 4, on Amazon, we can observe that the feature similarities between most benign nodes and their corresponding subgraphs are very close to 1.0, while fraud nodes are close to 0.0. This proves the effectiveness and superiority of using node-subgraph feature contrastive learning that integrates partially observed node labels to estimate the abnormal probability of nodes. Similarly, on YelpChi, while a few benign nodes have relatively low feature similarities with their corresponding subgraphs, the majority of benign nodes are above 0.8, or even infinitely close to 1.0. This also aligns with expectations, as for benign nodes, their embeddings may slightly deviate from those of the corresponding subgraphs due to differences in interests and hobbies. But anyway, in both these two datasets, our proposed node-subgraph contrastive strategy can well analyze whether nodes tend to be benigns or frauds from a feature perspective, thereby providing informative and useful guidance for dynamic neighborhood modeling.

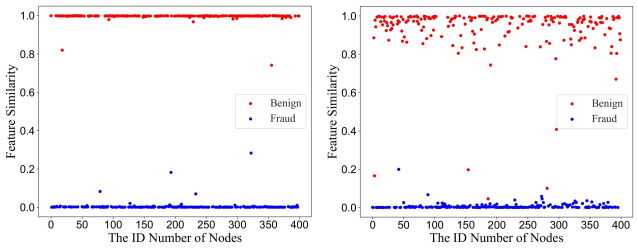

(a) Amazon (b) YelpChi

Figure 4: Analysis of feature similarity between embeddings of nodes and their corresponding subgraphs.

Parameter Sensitivity

We report Macro-F1 using the Amazon and YelpChi datasets as examples to investigate the sensitivity of two important hyperparameters, namely step size k and restart probability r for random walk with restart in subgraph sampling.

Analysis of Step Size k. The step size k determines the path length for random walk with restart. We vary its value from 4 to 12 with an increment of 2, and the results are presented in Figure 5. As shown, DCL-GFD achieves the best performance when k = 10 on Amazon and k = 12 on YelpChi, respectively. Moreover, the performance trends of these two datasets are generally consistent, that is, increasing first and then decreasing, reaching the optimal value at

a certain point. This is mainly due to the fact that a small step size k is not sufficient to obtain informative subgraph embeddings for estimate node abnormality, whereas a large step size k may introduce noise and weaken the modeling between the target node and its corresponding subgraphs.

Analysis of Restart Probability r. The restart probability r determines the likelihood of returning to the initial node after each walk. We vary its value from 0.1 to 0.9 and the corresponding results are reported in Figure 5. As shown, the performance of our DCL-GFD decreases in most cases when the restart probability r is too large or too small. This observation indicates that a small restart probability r leads to a tendency to move within local neighborhoods, inevitably ignoring important global information, whereas a large restart probability r lead to frequently returning to the target node and exploring extensively in the global scope, overlooking the specific local relationships among nodes.

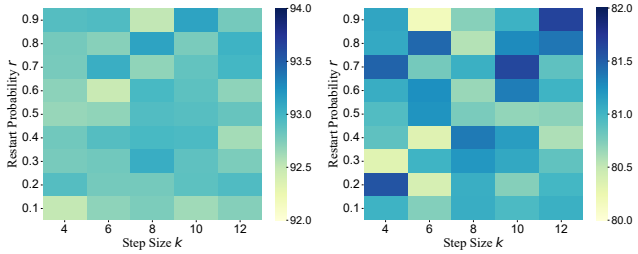

(a) Amazon (b) YelpChi

Figure 5: Parameter sensitivity analysis of k and r.

Related Work

In line with the focus of our work, we review some closely related studies, including graph neural networks, GNN-based fraud detection and graph contrastive learning.

Graph Neural Networks. Over the past decade, graph neural networks have been actively used in tackling graph-structured data. For instance, GCN (Kipf and Welling 2017) presents a representative graph convolutional operation that aggregates information from nodes' one-hop neighbor. GAT (Velickovic et al. 2018) assigns different weights to neighbors using a node-level attention mechanism. GraphSage (Hamilton, Ying, and Leskovec 2017) designs various pooling operations, including mean or max, to sample and aggregate information from neighbors. AM-GCN (Wang et al. 2020) designs a multi-channel GCN that adaptively learns node embeddings from both topology and feature space. More detailed surveys can be found in (Wu et al. 2021). Though these methods can be effectively used for learning node embeddings, the inherent class imbalance and heterophily issues within graph-based fraud detection scenarios pose significant challenges to existing GNNs.

GNN-Based Fraud Detection. Recently, much attention has been paid to introduce GNNs to identify frauds that deviate significantly from the majority of benigns (Zheng et al. 2024; Liu et al. 2020). For example, PC-GNN (Liu et al. 2021) designs a label-balance sampler which extracts nodes and edges to build training sub-graphs to remedy the class imbalance problem. GDN (Gao et al. 2023b) adopts different strategies to divide and constrain node features for different types of nodes, so as to solve the structural distribution shift from the perspective of feature. BSL (Yu, Liu, and Luo 2024) decouples node feature according to the type of edges, and performs strong and weak enhancement on unlabeled samples to assist in training. DGA-GNN (Duan et al. 2024) adopts decision tree binning encoding to convert non-additive node attributes into bin vectors, and designs a feedback dynamic grouping strategy to extract more discriminative features from neighbors. DiG-In-GNN (Zhang et al. 2024) employs contrastive learning to generate discriminative guide nodes and selects neighbors via reinforcement learning, thereby relieving the problem of feature and structural inconsistency caused by abnormal node camouflage. More detailed surveys can be found in (Tang et al. 2023). Nevertheless, few works model neighborhood by dynamically estimating the abnormal tendency of neighbor nodes during information propagation to enhance the effectiveness and robustness of fraud detection.

Graph Contrastive Learning. As one of the most effective self-supervised methods, contrastive learning (CL) aims to learn discriminative node embeddings by maximizing the mutual information of positive pairs while minimizing that of negative pairs. For instance, DGI (Velickovic et al. 2019) learns node embeddings by maximizing the mutual information between nodes and graphs. GraphCL (Hafidi et al. 2022) designs a new contrast strategy that maximizes the similarity between node embeddings from two randomly perturbed perspectives of the same graph. RGCL (Li et al. 2022) points that a high-performing augmentation should preserve the salient semantics of anchor graphs and create rationale-aware views for contrastive learning. HTML (Li et al. 2024) introduces knowledge distillation into the GCL to model topology level discriminative information. In general, our focus is mainly on how to use contrastive learning to facilitate graph-based fraud detection task.

Conclusion

In this paper, we aim to address the inherent issues of class imbalance and heterophily within graph-based fraud detection, and present a DCL-GFD approach for dynamic neighborhood modeling via node-subgraph contrastive learning. Specifically, we first estimate the abnormal probability of each node through node-subgraph contrast from the perspective of feature, which effectively maximizes the similarity between benigns and their corresponding subgraphs while minimizes that of frauds by optimizing contrastive learning objective with partially observed node labels. Based on the guidance of the probability that neighbors belonging to benigns or frauds, we then present a dynamic neighborhood modeling that learns benign-like and fraud-like neighbor embeddings, respectively. This essentially simulates the uncertain and mixed nature of neighbors associated with the target node by continuously adjusting their class tendency, thereby improving the discriminative ability of nodes. Extensive experiments across various public datasets demonstrate that DCL-GFD outperforms state-of-the-art baselines in terms of graph-based fraud detection.

Acknowledgments

This work was supported by the National Key Research and Development Program of China (No. 2023YFB2603904, No. 2023YFC3304503), the National Natural Science Foundation of China (No. 62402337, No. 92370111, No. 62422210, No. 62272340, No. 62276187), the Postdoctoral Fellowship Program of CPSF under Grant Number GZC20241207, and the China Postdoctoral Science Foundation under Grant Number 2024M752367.

References

Chai, Z.; You, S.; Yang, Y.; Pu, S.; Xu, J.; Cai, H.; and Jiang, W. 2022. Can Abnormality be Detected by Graph Neural Networks? In *Proceedings of the 31st International Joint Conference on Artificial Intelligence*, 1945–1951.

Chen, N.; Liu, Z.; Hooi, B.; He, B.; Fathony, R.; Hu, J.; and Chen, J. 2024. Consistency Training with Learnable Data Augmentation for Graph Anomaly Detection with Limited Supervision. In *Proceedings of the 12th International Conference on Learning Representations*.

Dong, Y.; He, D.; Wang, X.; Jin, Y.; Ge, M.; Yang, C.; and Jin, D. 2024. Unveiling Implicit Deceptive Patterns in Multi-Modal Fake News via Neuro-Symbolic Reasoning. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence*, 8354–8362.

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 and 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 37th AAAI Conference on Artificial Intelligence*, 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*.

Duan, M.; Zheng, T.; Gao, Y.; Wang, G.; Feng, Z.; and Wang, X. 2024. DGA-GNN: Dynamic Grouping Aggregation GNN for Fraud Detection. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence*, 11820–11828.

Gao, Y.; Fang, J.; Sui, Y.; Li, Y.; Wang, X.; Feng, H.; and Zhang, Y. 2024. Graph Anomaly Detection with Bi-level Optimization. In *Proceedings of the ACM Web Conference*, 4383–4394.

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*, 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 16th ACM International Conference on Web Search and Data Mining*, 357–365.

Hafidi, H.; Ghogho, M.; Ciblat, P.; and Swami, A. 2022. Negative Sampling Strategies for Contrastive Self-Supervised Learning of Graph Representations. *Signal Process*, 190: 108310.

Hamilton, W. L.; Ying, Z.; and Leskovec, J. 2017. Inductive Representation Learning on Large Graphs. In *Proceedings of Advances in Neural Information Processing Systems*, 1024–1034.

Jin, M.; Liu, Y.; Zheng, Y.; Chi, L.; Li, Y.; and Pan, S. 2021. ANEMONE: Graph Anomaly Detection with Multi-Scale Contrastive Learning. In *Proceedings of the 30th ACM International Conference on Information and Knowledge Management*, 3122–3126.

Kipf, T. N.; and Welling, M. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *Proceedings of the 5th International Conference on Learning Representations*.

Li, J.; Jin, Y.; Gao, H.; Qiang, W.; Zheng, C.; and Sun, F. 2024. Hierarchical Topology Isomorphism Expertise Embedded Graph Contrastive Learning. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence*, 13518–13527.

Li, S.; Wang, X.; Zhang, A.; Wu, Y.; He, X.; and Chua, T. 2022. Let Invariant Rationale Discovery Inspire Graph Contrastive Learning. In *Proceedings of International Conference on Machine Learning*, 13052–13065.

Liu, Y.; Ao, X.; Qin, Z.; Chi, J.; Feng, J.; Yang, H.; and He, Q. 2021. Pick and Choose: A GNN-based Imbalanced Learning Approach for Fraud Detection. In *Proceedings of the ACM Web Conference*, 3168–3177.

Liu, Z.; Dou, Y.; Yu, P. S.; Deng, Y.; and Peng, H. 2020. Alleviating the Inconsistency Problem of Applying Graph Neural Network to Fraud Detection. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 1569–1572.

McAuley, J. J.; and Leskovec, J. 2013. From Amateurs to Connoisseurs: Modeling the Evolution of User Expertise Through Online Reviews. In *Proceedings of the ACM Web Conference*, 897–908.

Qiu, J.; Chen, Q.; Dong, Y.; Zhang, J.; Yang, H.; Ding, M.; Wang, K.; and Tang, J. 2020. GCC: Graph Contrastive Coding for Graph Neural Network Pre-Training. In *Proceedings of the 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 1150–1160.

Rayana, S.; and Akoglu, L. 2015. Collective Opinion Spam Detection: Bridging Review Networks and Metadata. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 985–994.

Shi, F.; Cao, Y.; Shang, Y.; Zhou, Y.; Zhou, C.; and Wu, J. 2022. H2-FDetector: A GNN-based Fraud Detector with Homophilic and Heterophilic Connections. In *Proceedings of the ACM Web Conference*, 1486–1494.

Tang, J.; Hua, F.; Gao, Z.; Zhao, P.; and Li, J. 2023. GADBench: Revisiting and Benchmarking Supervised Graph Anomaly Detection. In *Proceedings of Advances in Neural Information Processing Systems*.

Tang, J.; Li, J.; Gao, Z.; and Li, J. 2022. Rethinking Graph Neural Networks for Anomaly Detection. In *Proceedings of International Conference on Machine Learning*, 21076–21089.

Velickovic, P.; Cucurull, G.; Casanova, A.; Romero, A.; Liò, P.; and Bengio, Y. 2018. Graph Attention Networks. In *Proceedings of the 6th International Conference on Learning Representations*.

Velickovic, P.; Fedus, W.; Hamilton, W. L.; Liò, P.; Bengio, Y.; and Hjelm, R. D. 2019. Deep Graph Infomax. In *Proceedings of the 7th International Conference on Learning Representations*.

Wang, M.; Yu, L.; Zheng, D.; Gan, Q.; Gai, Y.; Ye, Z.; Li, M.; Zhou, J.; Huang, Q.; and Ma, C. 2019. Deep Graph Library: Towards Efficient and Scalable Deep Learning on Graphs. In *Proceedings of ICLR Workshop on Representation Learning on Graphs and Manifolds*.

Wang, X.; Zhu, M.; Bo, D.; Cui, P.; Shi, C.; and Pei, J. 2020. AM-GCN: Adaptive Multi-channel Graph Convolutional Networks. In *Proceedings of the 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 1243–1253.

Wang, Y.; Zhang, J.; Huang, Z.; Li, W.; Feng, S.; Ma, Z.; Sun, Y.; Yu, D.; Dong, F.; Jin, J.; Wang, B.; and Luo, J. 2023. Label Information Enhanced Fraud Detection against Low Homophily in Graphs. In *Proceedings of the ACM Web Conference*, 406–416.

Wang, Z.; Mu, C.; Hu, S.; Chu, C.; and Li, X. 2022. Modelling the Dynamics of Regret Minimization in Large Agent Populations: a Master Equation Approach. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence*, 534–540.

Wu, B.; Yao, X.; Zhang, B.; Chao, K.; and Li, Y. 2023. SplitGNN: Spectral Graph Neural Network for Fraud Detection against Heterophily. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, 2737–2746.

Wu, Z.; Pan, S.; Chen, F.; Long, G.; Zhang, C.; and Yu, P. S. 2021. A Comprehensive Survey on Graph Neural Networks. *IEEE Transactions on Neural Networks and Learning Systems*, 32(1): 4–24.

Xu, F.; Wang, N.; Wu, H.; Wen, X.; Zhao, X.; and Wan, H. 2024. Revisiting Graph-Based Fraud Detection in Sight of Heterophily and Spectrum. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence*, 9214–9222.

You, Y.; Chen, T.; Sui, Y.; Chen, T.; Wang, Z.; and Shen, Y. 2020. Graph Contrastive Learning with Augmentations. In *Proceedings of Advances in Neural Information Processing Systems*.

Yu, H.; Liu, Z.; and Luo, X. 2024. Barely Supervised Learning for Graph-Based Fraud Detection. In *Proceedings of 38th AAAI Conference on Artificial Intelligence*, 16548–16557.

Yu, J.; Wang, H.; Wang, X.; Li, Z.; Qin, L.; Zhang, W.; Liao, J.; and Zhang, Y. 2023. Group-based Fraud Detection Network on e-Commerce Platforms. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 5463–5475.

Yun, S.; Kim, S.; Lee, J.; Kang, J.; and Kim, H. J. 2021. Neo-GNNs: Neighborhood Overlap-aware Graph Neural Networks for Link Prediction. In *Proceedings of Advances in Neural Information Processing Systems*, 13683–13694.

Zhang, G.; Wu, J.; Yang, J.; Beheshti, A.; Xue, S.; Zhou, C.; and Sheng, Q. Z. 2021. FRAUDRE: Fraud Detection Dual-Resistant to Graph Inconsistency and Imbalance. In *Proceedings of IEEE International Conference on Data Mining*, 867–876.

Zhang, J.; Xu, Z.; Lv, D.; Shi, Z.; Shen, D.; Jin, J.; and Dong, F. 2024. DiG-In-GNN: Discriminative Feature Guided GNN-Based Fraud Detector against Inconsistencies in Multi-Relation Fraud Graph. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence*, 9323–9331.

Zhao, T.; Deng, C.; Yu, K.; Jiang, T.; Wang, D.; and Jiang, M. 2020. Error-Bounded Graph Anomaly Loss for GNNs. In *Proceedings of the 29th ACM International Conference on Information and Knowledge Management*, 1873–1882.

Zheng, X.; Wu, B.; Liang, X.; and Li, W. 2024. Friend or Foe? Mining Suspicious Behavior via Graph Capsule Infomax Detector against Fraudsters. In *Proceedings of the ACM Web Conference*, 2684–2693.

Zhong, P.; Wang, D.; and Miao, C. 2019. An Affect-Rich Neural Conversational Model with Biased Attention and Weighted Cross-Entropy Loss. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence*, 7492–7500.

Zhu, J.; Yan, Y.; Zhao, L.; Heimann, M.; Akoglu, L.; and Koutra, D. 2020. Beyond Homophily in Graph Neural Networks: Current Limitations and Effective Designs. In *Proceedings of Advances in Neural Information Processing Systems*.

Zhuo, W.; Liu, Z.; Hooi, B.; He, B.; Tan, G.; Fathony, R.; and Chen, J. 2024. Partitioning Message Passing for Graph Fraud Detection. In *Proceedings of the 12th International Conference on Learning Representations*.