

Beyond Smoothing: Unsupervised Graph Representation Learning with Edge Heterophily Discriminating

Yixin Liu¹, Yizhen Zheng¹, Daokun Zhang¹, Vincent CS Lee¹, Shirui Pan²

¹ Monash University, Australia

² Griffith University, Australia

{yixin.liu, yizhen.zheng1, daokun.zhang, vincent.cs.lee}@monash.edu, s.pan@griffith.edu.au

Abstract

Unsupervised graph representation learning (UGRL) has drawn increasing research attention and achieved promising results in several graph analytic tasks. Relying on the homophily assumption, existing UGRL methods tend to smooth the learned node representations along all edges, ignoring the existence of heterophilic edges that connect nodes with distinct attributes. As a result, current methods are hard to generalize to heterophilic graphs where dissimilar nodes are widely connected, and also vulnerable to adversarial attacks. To address this issue, we propose a novel unsupervised **Graph Representation learning method with Edge Heterophily discriminating** (GREET) which learns representations by discriminating and leveraging homophilic edges and heterophilic edges. To distinguish two types of edges, we build an edge discriminator that infers edge homophily/heterophily from feature and structure information. We train the edge discriminator in an unsupervised way through minimizing the crafted pivot-anchored ranking loss, with randomly sampled node pairs acting as pivots. Node representations are learned through contrasting the dual-channel encodings obtained from the discriminated homophilic and heterophilic edges. With an effective interplaying scheme, edge discriminating and representation learning can mutually boost each other during the training phase. We conducted extensive experiments on 14 benchmark datasets and multiple learning scenarios to demonstrate the superiority of GREET.

Introduction

Unsupervised graph representation learning (UGRL) aims to learn low-dimensional representations from graph-structured data without costly labels (Kipf and Welling 2016; Hamilton, Ying, and Leskovec 2017; Veličković et al. 2019). Such representations can benefit myriads of downstream tasks, including node classification, link prediction and node clustering (Zhang, Yin, and Philip 2022; Pan et al. 2018). In recent years, UGRL has attracted increasing research attention and been applied to diverse applications, such as recommender systems (Ge et al. 2020), traffic prediction (Jenkins et al. 2019), and drug molecular modeling (Wang et al. 2021).

Most of the UGRL methods are designed based on the homophily assumption, i.e., linked nodes tend to share similar

attributes with each other (Hamilton 2020). Depending on such an assumption, they utilize low-pass filter-like graph neural networks (GNNs) (Wu et al. 2019; Kipf and Welling 2017; Li et al. 2019) to smooth the representations of adjacent nodes. Meanwhile, to provide supervision signals for representation learning, they often use objectives that preserve local smoothness, i.e., encouraging the representations of nodes within the same edge (Kipf and Welling 2016; Peng et al. 2020), random walk (Hamilton, Ying, and Leskovec 2017; Perozzi, Al-Rfou, and Skiena 2014), or subgraph (Qiu et al. 2020) to have higher similarity to each other. As a result, all the nodes are forced to have similar representations to their neighbors. As shown in Fig. 1(a), all the connected nodes are pushed to be closer in the representation space, even if some of them have moderate feature similarities that are comparable to randomly sampled node pairs.

With such a representation smoothing mechanism, existing UGRL methods cannot always generate optimal representations, since real-world graph data are likely to violate the homophily assumption (Zhu et al. 2020a; Bo et al. 2021). Actually, apart from *homophilic edges* that connect similar nodes, there widely exist *heterophilic edges* in real-world graphs, which otherwise link dissimilar nodes. Take social networks as an example, a person may have the same hobby with some friends, but also have quite different interests with other friends. Moreover, in a special type of graphs named heterophilic graphs (e.g., Web Knowledge Base graph Texas, as shown in Fig. 1(b)), heterophilic edges are far more than homophilic ones (Pei et al. 2020). Besides, graphs in reality inevitably contain noisy links caused by the uncertain factors in real-world systems or even adversarial attacks (Zhao et al. 2021a; Chang et al. 2021). In this situation, smoothing the information along all edges can lead to indistinguishable node representations, hindering the generalization ability and robustness of UGRL methods.

Considering the existence of homophilic and heterophilic edges, a natural question can be raised: (*Q1*) *Is it possible to distinguish between two types of edges in an unsupervised manner?* In supervised scenarios, edge discriminating can be designed as a built-in component in node classification models, with labels providing direct evidence (Zhao et al. 2021b). Nevertheless, in unsupervised scenarios, it is hard to discriminate edge types through only node features and graph structure, especially for the widely existing edges that

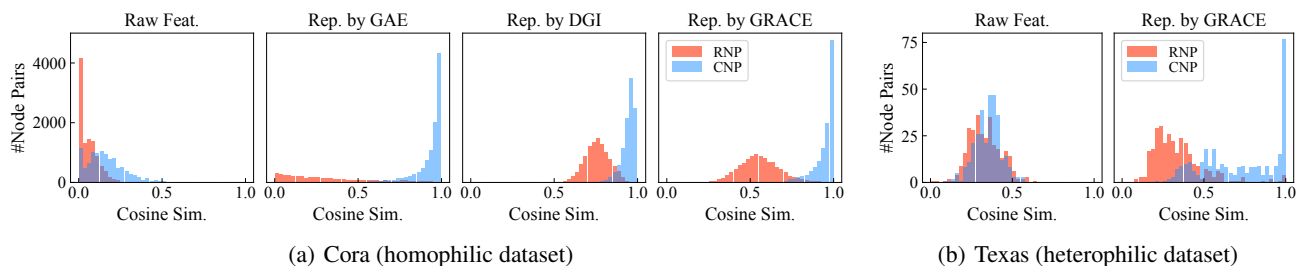


Figure 1: The distribution of pair-wise cosine similarity of connected node pairs (CNP) and randomly sampled node pairs (RNP). (a) The similarity w.r.t. raw features and representations learned by three UGRL methods (i.e., GAE (Kipf and Welling 2016), DGI (Veličković et al. 2019), and GRACE (Zhu et al. 2020b)) on Cora dataset. (b) The similarity w.r.t. raw features and representations learned by GRACE (Zhu et al. 2020b) on Texas dataset.

connect node pairs with moderate similarities. Even though a workable unsupervised edge discriminating strategy can be developed, we have to be faced with a follow-up challenge: *(Q2) How to effectively couple edge discriminating with representation learning into an integrated UGRL model?* On the one hand, for learning informative node representations, we have to make the best of the discriminated homophilic and heterophilic edges, which inevitably involve some errors. On the other hand, we shall find a good interplaying scheme to make edge discriminating and representation learning mutually boost each other.

To answer both of the aforementioned questions, in this paper, we propose a novel unsupervised **Graph Representation learning method with Edge hEterophilily discriminaTing** (GREET). Our theme is to devise an effective unsupervised edge discriminating component, and integrate it with a dual-channel graph encoding component into an end-to-end UGRL framework. More concretely, to answer *(Q1)*, for edge discriminating, we design a trainable edge discriminator with node features and structural encodings, and learn it in an unsupervised manner. To provide robust supervision, we design a pivot-anchored ranking loss, which uses the similarities of random node pairs as pivots, and evaluates the pivot-relative offside extent delivered by the end-node similarities of discriminated homophilic and heterophilic edges. Through minimizing the pivot-anchored ranking loss, we can effectively alleviate the misleading supervision caused by the ambiguous homophilic/heterophilic edges with moderate similarities. To answer *(Q2)*, with the discriminated homophilic and heterophilic edges, we design a dual-channel graph encoding component to generate node representations, with each channel playing an essential role on the corresponding edge type view. A robust cross-channel contrasting mechanism is developed to learn informative node representations. To make the representation learning well integrated with edge discriminating, we use the learned node representations to measure node similarities for training the edge discriminator. Performed with a closed-loop interplay, edge discriminating and representation learning can increasingly promote each other during the model training phase. We conduct extensive experiments on 14 benchmark datasets, and the results demonstrate the superior effectiveness and robustness of our proposed GREET over state-of-the-art methods.

Related Work

In this section, we briefly review two related research directions. A more detailed review is available in Appendix A.

Graph Neural Networks (GNNs) aim to model graph-structured data via graph convolution based on spectral theory (Bruna et al. 2014; Kipf and Welling 2017; Wu et al. 2019) or spatial information aggregation (Hamilton, Ying, and Leskovec 2017; Veličković et al. 2018; Xu et al. 2019). From the perspective of graph signal processing, most GNNs can be viewed as low-pass graph filters that smooth features over graph topology, resulting in similar representations between adjacent nodes (Nt and Maehara 2019; Wu et al. 2019). This property, unfortunately, hinders the capability of GNNs in heterophilic graphs where dissimilar nodes tend to be connected (Lim et al. 2021). Some recent GNNs try to tackle this issue with novel designs (Zheng et al. 2022a), such as advanced aggregation functions (Pei et al. 2020; Bo et al. 2021; Jin et al. 2021) and network architecture (Zhu et al. 2020a; Chien et al. 2021; Suresh et al. 2021). However, they mainly focus on semi-supervised learning, leaving works on unsupervised learning unexplored.

Unsupervised Graph Representation Learning (UGRL) aims to learn low-dimensional node representations on graphs. Early methods learn by maximizing the representation similarity among proximal nodes (within the same random walk (Perozzi, Al-Rfou, and Skiena 2014; Grover and Leskovec 2016) or same edge (Kipf and Welling 2016; Pan et al. 2018)). Recent efforts apply contrastive learning (Chen et al. 2020) to UGRL, which optimizes models by maximizing the agreement between two augmented graph views (Hassani and Khasahmadi 2020; Zhu et al. 2021; Veličković et al. 2019; Liu et al. 2022a,b, 2023). However, due to their low-pass filtering GNN encoders (Kipf and Welling 2017; Wu et al. 2019) and smoothness-increasing learning objectives (Perozzi, Al-Rfou, and Skiena 2014; Kipf and Welling 2016; Peng et al. 2020; Qiu et al. 2020), existing UGRL methods tend to smooth the representations of every connected node pair, leading to their sub-optimal performance, especially on noisy or heterophilic graphs. Another line of methods aims to learn representations from multiple edge-based pre-defined graph views (Qu et al. 2017; Park et al. 2020; Li et al. 2020b; Lee et al. 2020). Differently, GREET learns homophilic and heterophilic from graphs with a single edge type, which is more challenging.

Problem Definition

Notation. Define a graph as $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = \{v_1, \dots, v_n\}$ represents the node set and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ represents the edge set. The edge connecting node v_i and v_j is denoted as $e_{i,j}$. The numbers of nodes and edges are denoted as $|\mathcal{V}| = n$ and $|\mathcal{E}| = m$, respectively. Let $\mathbf{X} \in \mathbb{R}^{n \times d_f}$ denote the feature matrix, where the i -th row \mathbf{x}_i represents the d_f -dimensional feature vector of node v_i . We denote the adjacency matrix of \mathcal{G} as $\mathbf{A} \in \mathbb{R}^{n \times n}$, where $\mathbf{A}_{ij} = 1$ if $e_{i,j} \in \mathcal{E}$ and $\mathbf{A}_{ij} = 0$ otherwise. Using the feature matrix and adjacency matrix, the graph can also be written as $\mathcal{G} = (\mathbf{A}, \mathbf{X})$. The symmetric normalized adjacency matrix is denoted as $\tilde{\mathbf{A}} = \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}$, where \mathbf{D} is the diagonal degree matrix such that $\mathbf{D}_{ii} = \sum_j \mathbf{A}_{ij}$. The Laplacian matrix of the graph is defined as $\mathbf{L} = \mathbf{D} - \mathbf{A}$, and the symmetric normalized Laplacian matrix is $\tilde{\mathbf{L}} = \mathbf{I} - \tilde{\mathbf{A}}$.

Problem Description. We aim to solve the node-level unsupervised graph representation learning (UGRL) problem. The objective is to learn a representation mapping function $\mathcal{F} : \mathbb{R}^{n \times d_f} \times \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^{n \times d_r}$ to compute a representation matrix $\mathcal{F}(\mathbf{X}, \mathbf{A}) = \mathbf{H}$, where the i -th row represents \mathbf{h}_i , the low-dimensional (i.e., $d_r \ll d_f$) representation of node v_i . These representations can be saved and utilized for downstream tasks, such as node classification.

Methodology

This section details our proposed method termed GREET. As illustrated in Fig. 2, GREET is mainly composed of two components: *edge discriminating module* that discriminates the homophilic and heterophilic edges, and *dual-channel representation learning module* that leverages both types of edges to generate informative node representations. In *edge discriminating module*, the edges are separated into homophilic and heterophilic edges by an edge discriminator. In *dual-channel representation learning module*, node representations are constructed through applying a dual-channel encoding component to perform graph convolution with low-pass and high-pass filters on the two discriminated views respectively. The two main components are trained in a mutually boosting manner, with a pivot-anchored ranking loss used to train the edge discriminator, a robust dual-channel contrastive loss employed to learn informative node representations, and an effective alternating training strategy leveraged to make the training on the two components increasingly reinforce each other.

Edge Discriminating

In order to distinguish the homophilic and heterophilic edges, we construct an edge discriminator that learns to estimate the homophily probability of each edge. Given two adjacent nodes, their homophily probability is essentially related to two critical factors: *node features* which incorporate their semantic contents and *structural characteristics* which reflect their topological roles. A direct tool to integrate both structure and feature information is GNN (Kipf and Welling 2017); however, as we discussed before, GNN tends to smooth the features and then lose the heterophily-related knowledge. Thus, we consider an alternative solu-

tion, that is, modeling the structural characteristics with vectorial structural encoding (a.k.a. positional encoding) (Li et al. 2020a; Dwivedi et al. 2022; Tan et al. 2023). In this way, the effective evidence can be preserved through concatenating node structural encodings with raw features. Specifically, we employ a random walk diffusion process-based structural encoding (Dwivedi et al. 2022) to capture structural information. The d_s -dimensional structural encoding \mathbf{s}_i of node v_i can be computed by d_s -step random walk-based graph diffusion:

$$\mathbf{s}_i = \left[\mathbf{T}_{ii}, \mathbf{T}_{ii}^2, \dots, \mathbf{T}_{ii}^{d_s} \right] \in \mathbb{R}^{d_s}, \quad (1)$$

where $\mathbf{T} = \mathbf{A} \mathbf{D}^{-1}$ is the random walk transition matrix. Taking the raw features and structural encodings as input, the edge discriminator estimates the homophily probability with two MLP layers:

$$\begin{aligned} \mathbf{h}'_i &= \text{MLP}_1([\mathbf{x}_i \parallel \mathbf{s}_i]), \quad \mathbf{h}'_j = \text{MLP}_1([\mathbf{x}_j \parallel \mathbf{s}_j]), \\ \theta_{i,j} &= (\text{MLP}_2([\mathbf{h}'_i \parallel \mathbf{h}'_j]) + \text{MLP}_2([\mathbf{h}'_j \parallel \mathbf{h}'_i])) / 2, \end{aligned} \quad (2)$$

where $[\cdot \parallel \cdot]$ denotes the concatenation operation, \mathbf{h}'_i is an intermediate embedding of node v_i , and $\theta_{i,j}$ is the estimated homophily probability for $e_{i,j}$. To make the estimation not affected by the direction of edges, we apply the second MLP layer on the different orders of embedding concatenations.

With the estimated probability $\theta_{i,j}$, our purpose is to sample a binary homophily indicator $w_{i,j}$ for edge $e_{i,j}$ from the Bernoulli distribution $w_{i,j} \sim \text{Bernoulli}(\theta_{i,j})$, with $w_{i,j} = 1$ indicating homophily and $w_{i,j} = 0$ indicating heterophily. However, the sampling is non-differentiable, making the discriminator hard to train. To address this issue, we approximate the binary indicator $w_{i,j}$ with Gumbel-Max reparametrization trick (Jang, Gu, and Poole 2017; Maddison, Mnih, and Teh 2017). Specifically, the discrete homophily indicator $w_{i,j}$ is relaxed to a continuous homophily weight $\hat{w}_{i,j}$, which is computed by:

$$\hat{w}_{i,j} = \text{Sigmoid}\left(\frac{(\theta_{i,j} + \log \delta - \log(1 - \delta))}{\tau_g}\right), \quad (3)$$

where $\delta \sim \text{Uniform}(0, 1)$ is the sampled Gumbel random variate and $\tau_g > 0$ is the temperature hyper-parameter. $\hat{w}_{i,j}$ tends to be sharper (closer to 0 or 1) with τ_g closer to 0.

With the estimated edge homophily indicators, we discriminate the original graph $\mathcal{G} = (\mathbf{A}, \mathbf{X})$ into the homophilic and heterophilic graph views, $\mathcal{G}^{(hm)} = (\mathbf{A}^{(hm)}, \mathbf{X})$ and $\mathcal{G}^{(ht)} = (\mathbf{A}^{(ht)}, \mathbf{X})$. To make the edge discriminating trainable, instead of explicitly separating edges into homophilic and heterophilic categories, we assign a soft weight for each edge when it works on the homophilic/heterophilic view:

$$\mathbf{A}_{i,j}^{(hm)} = \hat{w}_{i,j}, \quad \mathbf{A}_{i,j}^{(ht)} = 1 - \hat{w}_{i,j}, \quad \text{for } e_{i,j} \in \mathcal{E}. \quad (4)$$

Dual-Channel Encoding

To learn informative representations from homophilic and heterophilic graph views, we design two different encoders that respectively perform low-pass and high-pass graph filtering on two distinct views. Our key motivation is to capture

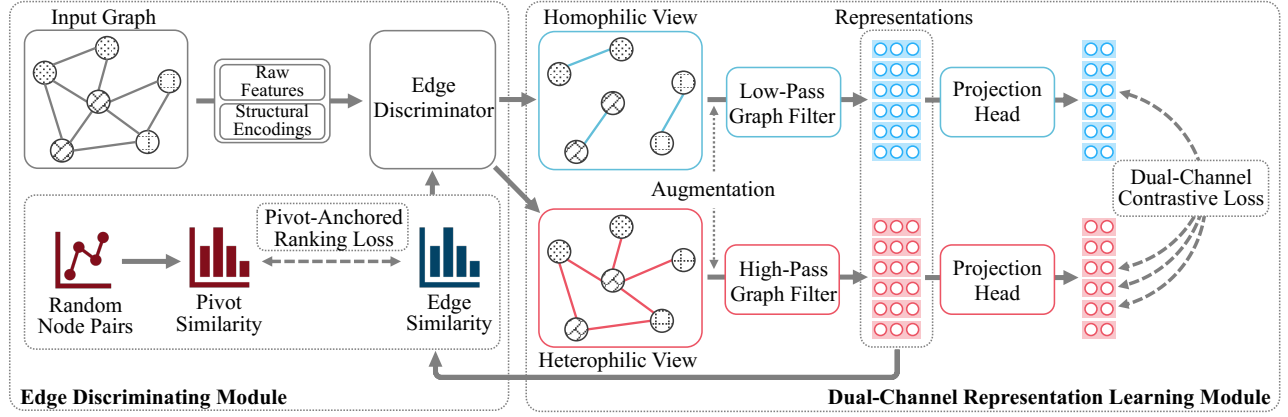


Figure 2: The overall pipeline of GREET. In the *edge discriminating module*, an edge discriminator is built to estimate the homophily probability of edges; in the *dual-channel representation learning module*, a dual-channel encoding component generates node representations from the discriminated homophilic and heterophilic edge views. The edge discriminator and encoding component are trained by minimizing a pivot-anchored ranking loss and a robust dual-channel contrastive loss, respectively.

the shared information among similar nodes according to homophilic edges, while filtering out irrelevant information from dissimilar neighbors according to heterophilic edges.

Concretely, on the homophilic view where similar nodes are connected to each other, we smooth the node features along the homophilic structure with a low-pass graph filter (Yang et al. 2021). The low-pass filtering can capture the low-frequency information in graph signals and then preserve the shared knowledge of similar nodes. In GREET, a simple low-pass GNN serves as the low-pass graph filter, which is expressed as:

$$\mathbf{H}_0^{(hm)} = \text{MLP}^{(hm)}(\mathbf{X}), \mathbf{H}_l^{(hm)} = \tilde{\mathbf{A}}^{(hm)} \mathbf{H}_{l-1}^{(hm)}, \quad (5)$$

where $\tilde{\mathbf{A}}^{(hm)}$ is the symmetric normalized homophilic adjacency matrix, $l \in \{1, \dots, L\}$ is the layer index, and the final output $\mathbf{H}_L^{(hm)}$ forms the homophilic representation matrix $\mathbf{H}^{(hm)}$.

Different from the homophilic view, the heterophilic view contains edges connecting dissimilar nodes. In this view, smoothing features with low-pass filtering would confuse the information of different nodes and lead to the loss of discriminative node properties (Bo et al. 2021). To better leverage the heterophilic edges, we use high-pass graph filtering, the inverse operation of low-pass filtering, to sharpen the node features along edges and preserve high-frequency graph signals. In this way, the representations of dissimilar but connected nodes can be distinguished from each other. From the perspective of graph signal processing, graph Laplacian operator (i.e., multiplying the graph Laplacian matrix \mathbf{L}) has been proven to be effective to capture high-frequency components (Dong et al. 2021; Ma et al. 2021). Accordingly, we design a simple GNN, namely Lap-GNN, to perform high-pass filtering on the heterophilic view:

$$\mathbf{H}_0^{(ht)} = \text{MLP}^{(ht)}(\mathbf{X}), \mathbf{H}_l^{(ht)} = \tilde{\mathbf{L}}^{(ht)} \mathbf{H}_{l-1}^{(ht)}, \quad (6)$$

where $l \in \{1, \dots, L\}$, and $\tilde{\mathbf{L}}^{(ht)} = \mathbf{I} - \alpha \tilde{\mathbf{A}}^{(ht)}$ is the symmetric normalized Laplacian matrix of heterophilic view

with α being a hyper-parameter to control the strength of high-pass filtering. We take the final output of Lap-GNN $\mathbf{H}_L^{(ht)}$ as the heterophilic representation matrix $\mathbf{H}^{(ht)}$.

Using two distinct encoders, now we acquire two groups of representations that capture low-frequency and high-frequency information respectively. Final node representations are obtained through concatenating the representations constructed from both views: $\mathbf{H} = [\mathbf{H}^{(hm)} \parallel \mathbf{H}^{(ht)}] \in \mathbb{R}^{n \times d_r}$, with its i -th row $\mathbf{h}_i \in \mathbb{R}^{d_r}$ being the representation vector of node v_i .

Model Training

To jointly learn edge distinction and node representations, we carefully design the learning objectives and training strategy for GREET. Specifically, we design a pivot-anchored ranking loss to train the edge discriminator, and construct a robust dual-channel contrastive loss to train the representation encoders. We also introduce an alternating training strategy to iteratively optimize two components.

Pivot-Anchored Ranking Loss. The target of the edge discriminator is to distinguish between homophilic edges (connecting similar nodes) and heterophilic edges (connecting dissimilar nodes), which is not a trivial task in an unsupervised scenario. The main challenge is to find the boundary between “similar” and “dissimilar”. To this end, we propose to use the randomly sampled node pairs as “pivots” for similarity measurement. Based on them, we introduce a pivot-anchored ranking loss to make sure that node pairs connected by homophilic edges are significantly more similar than the pivot node pairs, while the node pairs connected by heterophilic edges are markedly more different than the pivot node pairs. In detail, for each homophilic/heterophilic edge $e_{i,j} \in \mathcal{E}$, its homophilic pivot-anchored ranking loss $\mathcal{R}^{(hm)}(e_{i,j})$ and heterophilic pivot-anchored ranking loss $\mathcal{R}^{(ht)}(e_{i,j})$ are respectively defined as:

$$\begin{aligned} \mathcal{R}^{(hm)}(e_{i,j}) &= [s_{v_i, v_j'} - s_{e_{i,j}} + \gamma^{(hm)}]_+, \\ \mathcal{R}^{(ht)}(e_{i,j}) &= [s_{e_{i,j}} - s_{v_i, v_j'} + \gamma^{(ht)}]_+, \end{aligned} \quad (7)$$

where $[x]_+ = \max(x, 0)$ denotes only the positive case of x is considered; $s_{e_{i,j}} = \cos(\mathbf{h}_i, \mathbf{h}_j)$ and $s_{v_{i'}, v_{j'}} = \cos(\mathbf{h}_{i'}, \mathbf{h}_{j'})$ are respectively the representation-level cosine similarity between the node pair connected by edge $e_{i,j}$ and the cosine similarity between two randomly sampled nodes $v_{i'}$ and $v_{j'}$; $\gamma^{(hm)}$ and $\gamma^{(ht)}$ are respectively the margin parameters for the homophilic and heterophilic pivot-anchored ranking losses, which can force the similarity difference to reach a significant level.

Through performing expectation over all homophilic and heterophilic edges respectively, we can obtain the expected homophilic and heterophilic pivot-anchored ranking losses $\mathcal{L}_r^{(hm)}$ and $\mathcal{L}_r^{(ht)}$ as:

$$\begin{aligned}\mathcal{L}_r^{(hm)} &= \mathbb{E}_{e_{i,j} \sim \Pi(\hat{w}_{i,j})} \mathcal{R}^{(hm)}(e_{i,j}) = \frac{\sum_{e_{i,j} \in \mathcal{E}} \hat{w}_{i,j} \mathcal{R}^{(hm)}(e_{i,j})}{\hat{W}^{(hm)}}, \\ \mathcal{L}_r^{(ht)} &= \mathbb{E}_{e_{i,j} \sim \Pi(1-\hat{w}_{i,j})} \mathcal{R}^{(ht)}(e_{i,j}) = \frac{\sum_{e_{i,j} \in \mathcal{E}} (1-\hat{w}_{i,j}) \mathcal{R}^{(ht)}(e_{i,j})}{\hat{W}^{(ht)}},\end{aligned}\quad (8)$$

where $\Pi(\hat{w}_{i,j})$ and $\Pi(1-\hat{w}_{i,j})$ are the multinomial distributions parameterized by homophilic and heterophilic edge weights $\hat{w}_{i,j}$ and $1-\hat{w}_{i,j}$; $\hat{W}^{(hm)}$ and $\hat{W}^{(ht)}$ are the normalization terms which are respectively obtained by summing over all homophilic edge weights $\hat{w}_{i,j}$ and heterophilic edge weights $1-\hat{w}_{i,j}$. Through summing up the two expected losses, the overall pivot-anchored ranking loss to be minimized can be obtained as:

$$\mathcal{L}_r = \mathcal{L}_r^{(hm)} + \mathcal{L}_r^{(ht)}. \quad (9)$$

Robust Dual-Channel Contrastive Loss. To provide supervision for node representation learning, we use a robust contrasting mechanism that forces our model to generate semantically consistent representations from two distinct graph views. More importantly, such a learning target can effectively combat negative impacts caused by the small amounts of false discriminated edges, through implicitly using the generated representations to reconstruct feature-level similarity. Specially, for each node $v_i \in \mathcal{V}$, we first find another node v_j within v_i 's top- k similar node set measured by features, $\mathcal{N}_i = \text{kNN}(v_i, k)$, then compare their representations across homophilic and heterophilic views. To provide a fair similarity measure, we project node representations delivered by two views into a latent space with two learnable MLPs respectively, then use the similarity between the projections in the latent space as node representation similarity measuring metric. Inspired by the InfoNCE contrastive loss (Chen et al. 2020; Zhu et al. 2020b), the robust cross-channel contrastive loss is defined as:

$$\begin{aligned}\mathcal{L}_c &= -\frac{1}{n} \sum_{v_i \in \mathcal{V}} \left[\frac{1}{2|\mathcal{N}_i|} \sum_{v_j \in \mathcal{N}_i} \left(\log \frac{e^{\cos(\mathbf{z}_i^{(hm)}, \mathbf{z}_j^{(ht)})/\tau_c}}{\sum_{v_k \in \mathcal{V} \setminus v_i} e^{\cos(\mathbf{z}_i^{(hm)}, \mathbf{z}_k^{(ht)})/\tau_c}} \right. \right. \\ &\quad \left. \left. + \log \frac{e^{\cos(\mathbf{z}_i^{(ht)}, \mathbf{z}_j^{(hm)})/\tau_c}}{\sum_{v_k \in \mathcal{V} \setminus v_i} e^{\cos(\mathbf{z}_i^{(ht)}, \mathbf{z}_k^{(hm)})/\tau_c}} \right) \right],\end{aligned}\quad (10)$$

where $\mathbf{z}_i^{(hm)}$ and $\mathbf{z}_i^{(ht)}$ are respectively the projections of v_i 's homophilic-view and heterophilic-view representations

in the latent space, τ_c is a temperature hyper-parameter for contrastive learning and $\cos(\cdot, \cdot)$ is cosine similarity. In practice, we calculate contrastive loss in a mini-batch manner (Chen et al. 2020) for large graph datasets.

Training Strategy. In our method, the training of edge discriminating module relies on node representations to measure node similarity; the representation learning module, in turn, generates representations from the two views delivered by edge discriminating. To train two components effectively, we utilize an alternating training strategy to optimize edge discriminating and node representation learning alternately, with the two components boosting each other increasingly. To sum up, the overall optimization objective can be written as $\mathcal{L} = \mathcal{L}_r + \mathcal{L}_c$. To improve the generalization ability of the representation learning module, we adopt data augmentation that increases the diversity of data for model training (You et al. 2020). Concretely, we employ two simple but effective augmentation strategies, feature masking and edge dropping (Zhu et al. 2020b; You et al. 2020), to perturb the features and structures of both views. We present the algorithmic description in Appendix B.

Scalability Extension. To improve the scalability of GREET, we introduce the following mechanisms. (1) For structural encoding computation, we use sparse matrix multiplication to calculate graph diffusion, which avoids $\mathcal{O}(n^3)$ time complexity. (2) To find similar nodes in \mathcal{L}_c , we use locality-sensitive approximation algorithm (Fatemi, Asri, and Kazemi 2021) to obtain kNN nodes. (3) We calculate \mathcal{L}_c within a mini-batch of nodes instead of all nodes, reducing the time complexity from $\mathcal{O}(n^2)$ to $\mathcal{O}(nb)$, where b is the size of mini-batch. Note that the structural encodings and kNN nodes can be pre-computed at once before model training. Experiments on CoAuthor Physics dataset verify that GREET is able to learn on graphs with over 34k nodes and 495k edges. Detailed complexity analysis of GREET can be found in Appendix C.

Experiments

Experimental Settings

Datasets. We take transductive node classification as the downstream task to evaluate the effectiveness of the learned representations. Our experiments are conducted on 14 commonly used benchmark datasets, including 8 homophilic graph datasets (i.e., Cora, CiteSeer, PubMed, Wiki-CS, Amazon Computer, Amazon Photo, CoAuthor CS, and CoAuthor Physics (Sen et al. 2008; Mernyei and Cangea 2020; Shchur et al. 2018)) and 6 heterophilic graph datasets (i.e., Chameleon, Squirrel, Actor, Cornell, Texas, and Wisconsin (Pei et al. 2020)). We split all datasets following the public splits (Yang, Cohen, and Salakhudinov 2016; Kipf and Welling 2017; Pei et al. 2020) or commonly used splits (Zhu et al. 2021; Thakoor et al. 2022). The details of datasets are summarized in Appendix D.

Baselines. We compare GREET with three groups of baseline methods: (1) supervised/semi-supervised learning methods (i.e., GCN (Kipf and Welling 2017), GAT (Veličković et al. 2018), and MLP), (2) conventional UGRL methods (i.e., node2vec (Grover and Leskovec 2016), DeepWalk (Perozzi, Al-Rfou, and Skiena 2014),

Methods	Cora	CiteSeer	PubMed	Wiki-CS	Amz. Comp.	Amz. Photo	Co. CS	Co. Physics
GCN*	81.5	70.3	79.0	76.89±0.37	86.34±0.48	92.35±0.25	93.10±0.17	95.54±0.19
GAT*	83.0	72.5	79.0	77.42±0.19	87.06±0.35	92.64±0.42	92.41±0.27	95.45±0.17
MLP	56.11±0.34	56.91±0.42	71.35±0.05	72.02±0.21	73.88±0.10	78.54±0.05	90.42±0.08	93.54±0.05
DeepWalk	69.47±0.55	58.82±0.61	69.87±1.25	74.35±0.06	85.68±0.06	89.44±0.11	84.61±0.22	91.77±0.15
node2vec	71.24±0.89	47.64±0.77	66.47±1.00	71.79±0.05	84.39±0.08	89.67±0.12	85.08±0.03	91.19±0.04
GAE	71.07±0.39	65.22±0.43	71.73±0.92	70.15±0.01	85.27±0.19	91.62±0.13	90.01±0.71	94.92±0.07
VGAE	79.81±0.87	66.75±0.37	77.16±0.31	75.63±0.19	86.37±0.21	92.20±0.11	92.11±0.09	94.52±0.00
DGI	82.29±0.56	71.49±0.14	77.43±0.84	75.73±0.13	84.09±0.39	91.49±0.25	91.95±0.40	94.57±0.38
GMI	82.51±1.47	71.56±0.56	79.83±0.90	75.06±0.13	81.76±0.52	90.72±0.33	OOM	OOM
MVGRL	83.03±0.27	72.75±0.46	79.63±0.38	79.97±0.18	87.09±0.27	92.01±0.13	91.97±0.19	95.53±0.10
GRACE	80.08±0.53	71.41±0.38	80.15±0.34	79.16±0.36	87.21±0.44	92.65±0.32	92.78±0.23	95.39±0.32
GCA	80.39±0.42	71.21±0.24	80.37±0.75	79.35±0.12	87.84±0.27	92.78±0.17	93.32±0.12	95.87±0.15
BGRL	81.08±0.17	71.59±0.42	79.97±0.36	78.74±0.22	88.92±0.33	93.24±0.29	93.26±0.36	95.76±0.38
GREET	83.81±0.87	73.08±0.84	80.29±1.00	80.68±0.31	87.94±0.35	92.85±0.31	94.65±0.18	96.13±0.12

Table 1: Results in terms of classification accuracies (in percent \pm standard deviation) on homophilic benchmarks. “*” indicates that results are borrowed from the original papers. OOM indicates Out-Of-Memory on a 24GB GPU.

Methods	Cham.	Squ.	Actor	Cor.	Texas	Wis.
GCN	59.6	36.3	30.8	57.0	60.0	56.5
GAT	56.4	32.1	28.1	59.5	61.6	54.7
MLP	46.9	29.3	35.7	81.1	81.6	84.3
Geom*	60.9	38.1	31.6	60.8	67.6	64.1
H2GCN*	59.4	37.9	35.9	82.2	84.9	86.7
FAGCN	63.4	41.2	35.7	81.4	84.3	83.3
GPR	61.6	39.7	35.3	81.9	83.2	84.1
DW	47.7	32.9	22.8	39.2	46.5	33.5
n2v	41.9	22.8	28.3	42.9	41.9	37.5
GAE	33.8	28.0	28.0	58.9	58.6	52.6
VGAE	35.2	29.5	27.0	59.2	59.2	56.7
DGI	40.0	31.8	29.8	63.4	60.6	55.4
GMI	47.0	30.1	27.8	54.8	50.5	46.0
MVGRL	51.1	35.5	30.0	64.3	62.4	62.4
GRACE	48.1	31.3	29.0	54.9	57.6	50.0
GR-FA	52.7	36.0	32.6	67.6	64.1	63.7
GCA	49.8	35.5	29.7	55.4	59.5	50.8
BGRL	47.5	32.6	29.9	57.3	59.2	52.4
GREET	63.6	42.3	36.6	85.1	87.0	<u>84.9</u>

Table 2: Results on heterophilic benchmarks.

GAE, and VGAE (Kipf and Welling 2016)), and (3) contrastive UGRL methods (i.e., DGI (Veličković et al. 2019), GMI (Peng et al. 2020), MVGRL (Hassani and Khasahmadi 2020), GRACE (Zhu et al. 2020b), GCA (Zhu et al. 2021), and BGRL (Thakoor et al. 2022)). For the experiments on heterophilic graphs, we further consider 4 heterophily-aware GNNs as baselines, including Geom-GCN (Pei et al. 2020), H2GCN (Zhu et al. 2020a), FAGCN (Bo et al. 2021), and GPR-GNN (Chien et al. 2021). We further equip GRACE with FAGCN encoder (termed GR-FA) for comparison.

Experimental details. For GREET and the unsupervised baselines, we follow the evaluation protocol that is widely used in previous UGRL methods (Zhu et al. 2020b, 2021). Specifically, in training phase, the model is trained without supervision; then, in evaluation phase, the learned representations are frozen and used to train and test with a logistic re-

gression classifier. For the supervised/semi-supervised baselines, we directly train the models in an end-to-end fashion where labels are available for representation learning, and evaluate with the trained models. For all datasets, we report the averaged test accuracy and standard deviation over 10 runs of experiments. We conduct grid search to choose the best hyper-parameters on validation set. We also search for the best hyper-parameter for baseline methods during reproduction. Specific hyper-parameter settings and more implementation details are in Appendix E. The experiments that investigate the parameter sensitivity of GREET are demonstrated in Appendix F. The code of GREET is available at <https://github.com/yixinliu233/GREET>.

Experimental Results

Performance Comparison. The node classification results on 8 homophilic datasets and 6 heterophilic datasets are presented in Table 1 and Table 2, respectively. First of all, we observe that GREET outperforms all baseline methods in 10 out of 14 benchmarks and achieves the runner-up performance on the rest 4 benchmarks. The superior performance indicates the distinction and exploitation of homophilic and heterophilic edges can universally benefit representation learning, since two types of edges widely exist in diverse real-world graph data. In Table 2, we find that GREET significantly outperforms conventional and contrastive UGRL methods. The main reason is that these UGRL methods constantly smooth the representations along heterophilic edges, making the representations indistinguishable. In contrast, GREET detects the heterophilic edges and sharpens the representations along them with a high-pass graph filter, leading to even better performance than semi-supervised GNNs for heterophilic graphs. Notably, equipping contrastive UGRL methods with heterophily-aware encoders (e.g., GR-FA) only yields minor performance gain, indicating that adapting UGRL methods to heterophilic graphs needs crafted designs rather than simply modifying the encoder.

Robustness Analysis. In this experiment, we verify the robustness of GREET on graph data under adversarial attack.

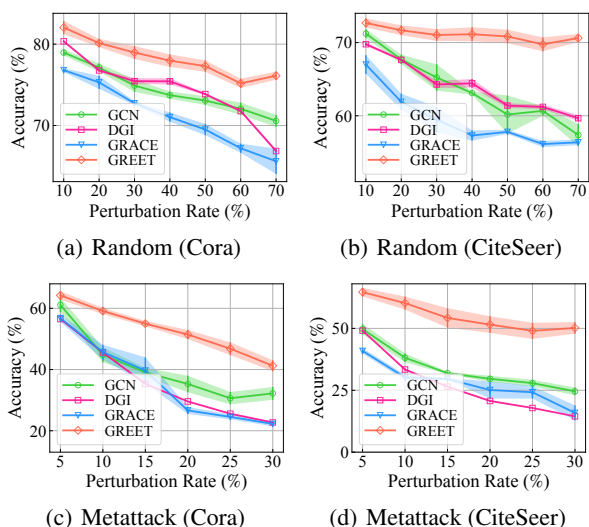


Figure 3: Performance under adversarial attacks.

Variants	Cornell	Texas	Cora	CiteSeer
GREET	85.14	87.03	83.81	73.08
w/o Eg. Dis.	82.16	82.97	82.19	72.18
w/o Du. Enc.	77.57	80.54	80.81	72.24
GNN Dis.	82.97	82.97	83.62	72.58
w/o Pivot	79.19	82.16	82.84	72.54
NCE Loss	77.57	80.00	82.48	70.00
Hom. Rep.	75.95	67.57	83.16	70.60
Het. Rep.	82.97	83.78	66.77	68.76

Table 3: Performance of GREET and its variants.

We perturb the graph structure with two non-targeted adversarial attack methods, i.e., random attack (randomly adding noisy edges to the original topology) and Metattack (Zügner and Günnemann 2019), and test the node classification accuracy on the representations learned from perturbed graphs. The node classification results under different perturbation rates are shown in Fig. 3. From the figure, we can witness that GREET consistently outperforms the baselines by a large margin. Moreover, with the increase of perturbation rate, the advantage of our method becomes more significant. The experimental results demonstrate the strong robustness of GREET against adversarial attacks on graph structures.

Ablation Study. To examine the contributions of each component and key design in GREET, we conduct experiments on several variants of GREET and the results are shown in Table 3. We first remove the key components, i.e. edge discriminator (w/o Eg. Dis.) and dual-channel encoding (w/o Du. Enc.), to investigate their effects. The results show that both components are critical to GREET, and the dual-channel encoding seems to contribute more. Then, we discuss the effects of three key designs (i.e., structural encoding in Eq.(2), pivot-anchored ranking loss, and robust dual-channel contrastive loss) by replacing them with alternative designs (i.e., GNN-based discriminator without structural encodings, ranking loss between homophilic and heterophilic edges without pivot, and InfoNCE (Chen




Figure 4: The distribution of pair-wise cosine similarity of connected node pairs (CNP) and randomly sampled node pairs (RNP) w.r.t. the representations learned by GREET on (a) Cora dataset, (b) Cora dataset under random attack with 50% perturbation rate, and (c) Texas dataset.

et al. 2020) loss) respectively. We can witness that GREET consistently outperforms three variants, which demonstrates the superiority of our designs. Finally, we investigate the quality of the representations generated from homophily view (Hom. Rep.) and heterophilic view (Het. Rep.). As expected, the homophilic representations are more effective on homophilic graphs (i.e., Cora and CiteSeer), while the heterophilic representations favor heterophilic graphs (i.e., Cornell and Texas). Jointly considering the representations from two views can produce the best results, because both of them contain critical and distinctive information from different perspectives.

Visualization. To investigate the property of the representations learned by GREET, we visualize the pair-wise representation similarity of Cora dataset, perturbed Cora dataset, and Texas dataset. As shown in Fig. 4(a), for most edges in Cora dataset, the representations of end nodes are similar; notably, for a small fraction of edges identified to be heterophilic, GREET forces their similarities to be close to 0. Such observation proves that GREET can separate homophilic and heterophilic edges and generate distinguishable representations accordingly. In the perturbed Cora dataset (Fig. 4(b)), a considerable part of edges is detected to be heterophilic and then have dissimilar end node representations. Thanks to such a capability of identifying the noisy edges, GREET shows strong robustness against structural attack in Fig. 3(a). A similar phenomenon can be found in the heterophilic dataset Texas (Fig. 4(c)) where the bulk of edges are recognized as heterophilic edges. The separation of two types of edges brings more informative representations, resulting in the superior performance of GREET on heterophilic graphs.

Conclusion

In this paper, we propose a novel method named GREET for unsupervised graph representation learning (UGRL). Our core idea is to discriminate and leverage homophilic and heterophilic edges to generate high-quality node representations. We employ an edge discriminator to distinguish two types of edges, and construct a dual-channel encoding component to generate representations according to edge distinction. We carefully design a pivot-anchored ranking loss and robust dual-channel contrastive loss for model training, with an alternating strategy to train both components in a mutually boosting manner. Extensive experiments reveal the effectiveness and robustness of our method.

Acknowledgements

The corresponding author is Shirui Pan. This work was supported by ARC Future Fellowship (No. FT210100097).

References

- Bo, D.; Wang, X.; Shi, C.; and Shen, H. 2021. Beyond Low-frequency Information in Graph Convolutional Networks. In *AAAI*, 3950–3957.
- Bruna, J.; Zaremba, W.; Szlam, A.; and LeCun, Y. 2014. Spectral networks and locally connected networks on graphs. In *ICLR*.
- Chang, H.; Rong, Y.; Xu, T.; Bian, Y.; Zhou, S.; Wang, X.; Huang, J.; and Zhu, W. 2021. Not All Low-Pass Filters are Robust in Graph Convolutional Networks. In *NeurIPS*.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A simple framework for contrastive learning of visual representations. In *ICML*, 1597–1607. PMLR.
- Chien, E.; Peng, J.; Li, P.; and Milenkovic, O. 2021. Adaptive Universal Generalized PageRank Graph Neural Network. In *ICLR*.
- Defferrard, M.; Bresson, X.; and Vandergheynst, P. 2016. Convolutional neural networks on graphs with fast localized spectral filtering. In *NeurIPS*, volume 29.
- Ding, K.; Xu, Z.; Tong, H.; and Liu, H. 2022. Data augmentation for deep graph learning: A survey. *SIGKDD Explorations*.
- Dong, Y.; Ding, K.; Jalaian, B.; Ji, S.; and Li, J. 2021. AdaGNN: Graph Neural Networks with Adaptive Frequency Response Filter. In *CIKM*, 392–401.
- Dwivedi, V. P.; Luu, A. T.; Laurent, T.; Bengio, Y.; and Bresson, X. 2022. Graph Neural Networks with Learnable Structural and Positional Representations. In *ICLR*.
- Fatemi, B.; Asri, L. E.; and Kazemi, S. M. 2021. SLAPS: Self-Supervision Improves Structure Learning for Graph Neural Networks. In *NeurIPS*.
- Fu, D.; Xu, Z.; Li, B.; Tong, H.; and He, J. 2020. A view-adversarial framework for multi-view network embedding. In *CIKM*, 2025–2028.
- Ge, S.; Wu, C.; Wu, F.; Qi, T.; and Huang, Y. 2020. Graph enhanced representation learning for news recommendation. In *WWW*, 2863–2869.
- Grover, A.; and Leskovec, J. 2016. node2vec: Scalable feature learning for networks. In *SIGKDD*, 855–864.
- Hamilton, W.; Ying, Z.; and Leskovec, J. 2017. Inductive representation learning on large graphs. In *NeurIPS*.
- Hamilton, W. L. 2020. Graph representation learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 14(3): 1–159.
- Hassani, K.; and Khasahmadi, A. H. 2020. Contrastive multi-view representation learning on graphs. In *ICML*, 4116–4126. PMLR.
- He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. 2020. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 9729–9738.
- Jang, E.; Gu, S.; and Poole, B. 2017. Categorical reparameterization with gumbel-softmax. In *ICLR*.
- Jenkins, P.; Farag, A.; Wang, S.; and Li, Z. 2019. Unsupervised representation learning of spatial data via multimodal embedding. In *CIKM*, 1993–2002.
- Jin, D.; Yu, Z.; Huo, C.; Wang, R.; Wang, X.; He, D.; and Han, J. 2021. Universal Graph Convolutional Networks. In *NeurIPS*, volume 34.
- Kipf, T. N.; and Welling, M. 2016. Variational graph auto-encoders. In *NeurIPS Workshop*.
- Kipf, T. N.; and Welling, M. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *ICLR*.
- Lee, Y.-C.; Seo, N.; Han, K.; and Kim, S.-W. 2020. Asine: Adversarial signed network embedding. In *SIGIR*, 609–618.
- Li, P.; Wang, Y.; Wang, H.; and Leskovec, J. 2020a. Distance encoding: Design provably more powerful neural networks for graph representation learning. In *NeurIPS*, volume 33.
- Li, Q.; Wu, X.-M.; Liu, H.; Zhang, X.; and Guan, Z. 2019. Label efficient semi-supervised learning via graph filtering. In *CVPR*, 9582–9591.
- Li, Y.; Tian, Y.; Zhang, J.; and Chang, Y. 2020b. Learning signed network embedding via graph attention. In *AAAI*, volume 34, 4772–4779.
- Lim, D.; Hohne, F.; Li, X.; Huang, S. L.; Gupta, V.; Bhalerao, O.; and Lim, S. N. 2021. Large Scale Learning on Non-Homophilous Graphs: New Benchmarks and Strong Simple Methods. In *NeurIPS*, volume 34.
- Liu, Y.; Ding, K.; Liu, H.; and Pan, S. 2023. GOOD-D: On Unsupervised Graph Out-Of-Distribution Detection. In *WSDM*.
- Liu, Y.; Jin, M.; Pan, S.; Zhou, C.; Zheng, Y.; Xia, F.; and Yu, P. 2022a. Graph self-supervised learning: A survey. *IEEE TKDE*.
- Liu, Y.; Li, Z.; Pan, S.; Gong, C.; Zhou, C.; and Karypis, G. 2021. Anomaly detection on attributed networks via contrastive self-supervised learning. *IEEE TNNLS*.
- Liu, Y.; Zheng, Y.; Zhang, D.; Chen, H.; Peng, H.; and Pan, S. 2022b. Towards unsupervised deep graph structure learning. In *WWW*, 1392–1403.
- Ma, Y.; Liu, X.; Zhao, T.; Liu, Y.; Tang, J.; and Shah, N. 2021. A unified view on graph neural networks as graph signal denoising. In *CIKM*, 1202–1211.
- Ma, Y.; Wang, S.; Aggarwal, C. C.; Yin, D.; and Tang, J. 2019. Multi-dimensional graph convolutional networks. In *SDM*, 657–665. SIAM.
- Maddison, C. J.; Mnih, A.; and Teh, Y. W. 2017. The concrete distribution: A continuous relaxation of discrete random variables. In *ICLR*.
- Mernyei, P.; and Cangea, C. 2020. Wiki-cs: A wikipedia-based benchmark for graph neural networks. *arXiv preprint arXiv:2007.02901*.
- Ni, J.; Chang, S.; Liu, X.; Cheng, W.; Chen, H.; Xu, D.; and Zhang, X. 2018. Co-regularized deep multi-network embedding. In *WWW*, 469–478.

- Nt, H.; and Maehara, T. 2019. Revisiting graph neural networks: All we have is low-pass filters. *arXiv preprint arXiv:1905.09550*.
- Pan, S.; Hu, R.; Long, G.; Jiang, J.; Yao, L.; and Zhang, C. 2018. Adversarially regularized graph autoencoder for graph embedding. In *IJCAI*, 2609–2615.
- Park, C.; Kim, D.; Han, J.; and Yu, H. 2020. Unsupervised attributed multiplex network embedding. In *AAAI*, volume 34, 5371–5378.
- Pei, H.; Wei, B.; Chang, K. C.-C.; Lei, Y.; and Yang, B. 2020. Geom-GCN: Geometric Graph Convolutional Networks. In *ICLR*.
- Peng, Z.; Huang, W.; Luo, M.; Zheng, Q.; Rong, Y.; Xu, T.; and Huang, J. 2020. Graph representation learning via graphical mutual information maximization. In *WWW*.
- Perozzi, B.; Al-Rfou, R.; and Skiena, S. 2014. Deepwalk: Online learning of social representations. In *SIGKDD*.
- 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 *SIGKDD*.
- Qu, M.; Tang, J.; Shang, J.; Ren, X.; Zhang, M.; and Han, J. 2017. An attention-based collaboration framework for multi-view network representation learning. In *CIKM*.
- Sen, P.; Namata, G.; Bilgic, M.; Getoor, L.; Galligher, B.; and Eliassi-Rad, T. 2008. Collective classification in network data. *AI magazine*, 29(3): 93–93.
- Shchur, O.; Mumme, M.; Bojchevski, A.; and Günnemann, S. 2018. Pitfalls of graph neural network evaluation. In *NeurIPS Workshop*.
- Sun, Y.; Bui, N.; Hsieh, T.-Y.; and Honavar, V. 2018. Multi-view network embedding via graph factorization clustering and co-regularized multi-view agreement. In *ICDM Workshop*, 1006–1013. IEEE.
- Suresh, S.; Budde, V.; Neville, J.; Li, P.; and Ma, J. 2021. Breaking the Limit of Graph Neural Networks by Improving the Assortativity of Graphs with Local Mixing Patterns. In *SIGKDD*, 1541–1551.
- Tan, Y.; Liu, Y.; Long, G.; Jiang, J.; Lu, Q.; and Zhang, C. 2023. Federated Learning on Non-IID Graphs via Structural Knowledge Sharing. In *AAAI*.
- Tan, Y.; Long, G.; Liu, L.; Zhou, T.; Lu, Q.; Jiang, J.; and Zhang, C. 2022a. Fedproto: Federated prototype learning across heterogeneous clients. In *AAAI*, volume 1, 3.
- Tan, Y.; Long, G.; Ma, J.; Liu, L.; Zhou, T.; and Jiang, J. 2022b. Federated Learning from Pre-Trained Models: A Contrastive Learning Approach. In *NeurIPS*.
- Tang, J.; Qu, M.; Wang, M.; Zhang, M.; Yan, J.; and Mei, Q. 2015. Line: Large-scale information network embedding. In *WWW*, 1067–1077.
- Thakoor, S.; Tallec, C.; Azar, M. G.; Azabou, M.; Dyer, E. L.; Munos, R.; Veličković, P.; and Valko, M. 2022. Large-scale representation learning on graphs via bootstrapping. In *ICLR*.
- Veličković, P.; Fedus, W.; Hamilton, W. L.; Liò, P.; Bengio, Y.; and Hjelm, R. D. 2019. Deep Graph Infomax. In *ICLR*.
- Veličković, P.; Cucurull, G.; Casanova, A.; Romero, A.; Liò, P.; and Bengio, Y. 2018. Graph Attention Networks. In *ICLR*.
- Wang, Y.; Min, Y.; Chen, X.; and Wu, J. 2021. Multi-view graph contrastive representation learning for drug-drug interaction prediction. In *WWW*, 2921–2933.
- Wu, F.; Souza, A.; Zhang, T.; Fifty, C.; Yu, T.; and Weinberger, K. 2019. Simplifying graph convolutional networks. In *ICML*, 6861–6871. PMLR.
- Wu, Z.; Pan, S.; Chen, F.; Long, G.; Zhang, C.; and Philip, S. Y. 2020. A comprehensive survey on graph neural networks. *IEEE TNNLS*, 32(1): 4–24.
- Xu, K.; Hu, W.; Leskovec, J.; and Jegelka, S. 2019. How Powerful are Graph Neural Networks? In *ICLR*.
- Yang, L.; Li, M.; Liu, L.; Wang, C.; Cao, X.; Guo, Y.; et al. 2021. Diverse Message Passing for Attribute with Heterophily. In *NeurIPS*, volume 34.
- Yang, Z.; Cohen, W.; and Salakhudinov, R. 2016. Revisiting semi-supervised learning with graph embeddings. In *ICML*, 40–48. PMLR.
- You, Y.; Chen, T.; Sui, Y.; Chen, T.; Wang, Z.; and Shen, Y. 2020. Graph contrastive learning with augmentations. In *NeurIPS*, volume 33, 5812–5823.
- Zhang, D.; Yin, J.; and Philip, S. Y. 2022. Link Prediction with Contextualized Self-Supervision. *IEEE TKDE*.
- Zhang, H.; Wu, B.; Yuan, X.; Pan, S.; Tong, H.; and Pei, J. 2022. Trustworthy Graph Neural Networks: Aspects, Methods and Trends. *arXiv preprint arXiv:2205.07424*.
- Zhao, J.; Wang, X.; Shi, C.; Hu, B.; Song, G.; and Ye, Y. 2021a. Heterogeneous graph structure learning for graph neural networks. In *AAAI*.
- Zhao, T.; Liu, Y.; Neves, L.; Woodford, O.; Jiang, M.; and Shah, N. 2021b. Data Augmentation for Graph Neural Networks. In *AAAI*, volume 35, 11015–11023.
- Zheng, X.; Liu, Y.; Pan, S.; Zhang, M.; Jin, D.; and Yu, P. S. 2022a. Graph neural networks for graphs with heterophily: A survey. *arXiv preprint arXiv:2202.07082*.
- Zheng, Y.; Pan, S.; Lee, V. C.; Zheng, Y.; and Yu, P. S. 2022b. Rethinking and Scaling Up Graph Contrastive Learning: An Extremely Efficient Approach with Group Discrimination. In *NeurIPS*.
- Zheng, Y.; Zheng, Y.; Zhou, X.; Gong, C.; Lee, V.; and Pan, S. 2022c. Unifying Graph Contrastive Learning with Flexible Contextual Scopes. In *ICDM*.
- Zhu, J.; Yan, Y.; Zhao, L.; Heimann, M.; Akoglu, L.; and Koutra, D. 2020a. Beyond Homophily in Graph Neural Networks: Current Limitations and Effective Designs. In *NeurIPS*, volume 33.
- Zhu, Y.; Xu, Y.; Yu, F.; Liu, Q.; Wu, S.; and Wang, L. 2020b. Deep Graph Contrastive Representation Learning. In *ICML Workshop*.
- Zhu, Y.; Xu, Y.; Yu, F.; Liu, Q.; Wu, S.; and Wang, L. 2021. Graph contrastive learning with adaptive augmentation. In *WWW*, 2069–2080.
- Zügner, D.; and Günnemann, S. 2019. Adversarial Attacks on Graph Neural Networks via Meta Learning. In *ICLR*.