

Restructuring Graphs for Higher Homophily via Adaptive Spectral Clustering

Shouheng Li^{1,3}, Dongwoo Kim², Qing Wang¹

¹ School of Computing, Australian National University, Canberra, Australia

² CSE & GSAI, POSTECH, Pohang, South Korea

³ Data61, CSIRO, Canberra, Australia

shouheng.li@anu.edu.au, dongwoo.kim@postech.ac.kr, qing.wang@anu.edu.au

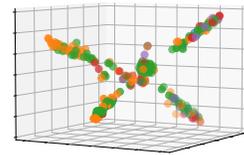
Abstract

While a growing body of literature has been studying new Graph Neural Networks (GNNs) that work on both homophilic and heterophilic graphs, little has been done on adapting classical GNNs to less-homophilic graphs. Although the ability to handle less-homophilic graphs is restricted, classical GNNs still stand out in several nice properties such as efficiency, simplicity, and explainability. In this work, we propose a novel graph restructuring method that can be integrated into any type of GNNs, including classical GNNs, to leverage the benefits of existing GNNs while alleviating their limitations. Our contribution is threefold: *a*) learning the weight of *pseudo-eigenvectors* for an adaptive spectral clustering that aligns well with known node labels, *b*) proposing a new density-aware homophilic metric that is robust to label imbalance, and *c*) reconstructing the adjacency matrix based on the result of adaptive spectral clustering to maximize homophilic scores. The experimental results show that our graph restructuring method can significantly boost the performance of six classical GNNs by an average of 25% on less-homophilic graphs. The boosted performance is comparable to state-of-the-art methods.

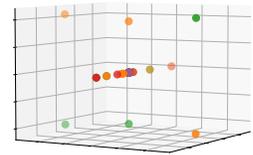
1 Introduction

Graph Neural Networks (GNNs) were originally inspired under the homophilic assumption - nodes of the same label are more likely to be connected than nodes of different labels. Recent studies have revealed the limitations of this homophilic assumption when applying GNNs on less-homophilic or heterophilic graphs (Pei et al. 2020). Since then, a number of approaches have been proposed with a focus on developing deep learning architectures for heterophilic graphs (Zhu et al. 2021b; Kim and Oh 2021; Chien et al. 2021; Li, Kim, and Wang 2021; Bo et al. 2021; Lim et al. 2021). However, little work has been done on adapting existing GNNs to less-homophilic graphs. Although existing GNNs such as GCN (Kipf and Welling 2017), ChevNet (Defferrard, Bresson, and Vandergheynst 2016), and GAT (Velickovic et al. 2018) are lacking the ability to work with less-homophilic graphs, they still stand out in several nice properties such as efficiency (Zeng et al. 2020), simplicity (Wu et al. 2019), and explainability (Ying et al. 2019). Our work aims to develop a graph restructuring method that can restructure a graph to

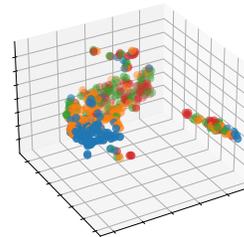
Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.



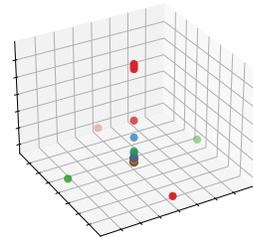
(a) WISCONSIN: visualized using T-SNE with the leading 5 eigenvectors



(b) WISCONSIN: visualized using T-SNE with the 22th, 44th and 206th eigenvectors



(c) EUROPE AIRPORT: visualized using T-SNE with the leading 5 eigenvectors



(d) EUROPE AIRPORT: visualized using the 366th, 382th and 3rd eigenvectors

Figure 1: Node clusters using different eigenvector choices on WISCONSIN and EUROPE AIRPORT. Colours represent node labels. Coordinates in 1a and 1c are computed using three-dimensional T-SNE. In 1b and 1d many nodes are overlapped so they appear to have fewer nodes than they actually have.

leverage the benefit of prevalent homophilic GNNs. Inspired by Klicpera, Weißenberger, and Günnemann (2019), we extend the restructuring approach to heterophilic graphs and boost the performance of GNNs that do not work well on less-homophilic graphs.

Discrepancy between SC and labels Spectral clustering (SC) aims to cluster nodes such that the edges between different clusters have low weights and the edges within a cluster have high weights. Such clusters are likely to align with node labels where a graph is homophilic, i.e., nodes with the

same labels are likely to connect closely. Nonetheless, this may not hold for less-homophilic graphs. Figure 1 visualizes the nodes in WISCONSIN (Pei et al. 2020) and EUROPE AIRPORT (Ribeiro, Saverese, and Figueiredo 2017) datasets based on the eigenvectors corresponding to the five smallest eigenvalues. As shown in Figure 1a and 1c, the labels are not aligned with the node clusters. However, if we choose the eigenvectors carefully, nodes clusters can align better with their labels, as evidenced by Figure 1b and 1d, which are visualized using manually chosen eigenvectors.

This observation shows that the eigenvectors corresponding to the leading eigenvalues do not always align well with the node labels. Particularly, in a heterophilic graph, two adjacent nodes are unlikely to have the same label, which is in contradiction with the smoothness properties of leading eigenvectors. However, when we choose eigenvectors appropriately, the correlation between the similarity of spectral features and node labels increases. To generalize this observation, we propose an adaptive spectral clustering algorithm that, *a*) divides the Laplacian spectrum into even slices, each slice corresponding to an embedding matrix called *pseudo-eigenvector*; *b*) learns the weights of *pseudo-eigenvectors* from existing node labels; *c*) restructures the graph according to node embedding distance to maximize homophily. To measure the homophilic level of graphs, we further introduce a new density-aware metric that is robust to label imbalance and label numbers. Our experimental results show that the performances of node-level prediction tasks with restructured graphs are greatly improved on classical GNNs.

2 Background

Spectral filtering. Let $\mathcal{G} = (V, E, \mathbf{A}, \mathbf{X})$ be an undirected graph with N nodes, where V , E , and $\mathbf{A} \in \{0, 1\}^{N \times N}$ are the node set, edge set, and adjacency matrix of \mathcal{G} , respectively, and $\mathbf{X} \in \mathbb{R}^{N \times F}$ is the node feature matrix. The normalized Laplacian matrix of \mathcal{G} is defined as $\mathbf{L} = \mathbf{I} - \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}$, where \mathbf{I} is the identity matrix with N diagonal entries and $\mathbf{D} \in \mathbb{R}^{N \times N}$ is the diagonal degree matrix of \mathcal{G} . In spectral graph theory, the eigenvalues $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_N)$ and eigenvectors \mathbf{U} of $\mathbf{L} = \mathbf{U} \Lambda \mathbf{U}^H$ are known as the graph’s spectrum and spectral basis, respectively, where \mathbf{U}^H is the Hermitian transpose of \mathbf{U} . The graph Fourier transform takes the form of $\hat{\mathbf{X}} = \mathbf{U}^H \mathbf{X}$ and its inverse is $\mathbf{X} = \mathbf{U} \hat{\mathbf{X}}$.

It is known that the Laplacian spectrum and spectral basis carry important information on the connectivity of a graph (Shuman et al. 2013). Lower frequencies correspond to global and smooth information on a graph, while higher frequencies correspond to local information, discontinuities and possible noise (Shuman et al. 2013). One can apply a spectral filter and use graph Fourier transform to manipulate signals on a graph in various ways, such as smoothing and denoising (Schaub and Segarra 2018), abnormally detection (Miller, Beard, and Bliss 2011) and clustering (Wai et al. 2018). Spectral convolution on graphs is defined as the multiplication of a signal \mathbf{x} with a filter $g(\Lambda)$ in the Fourier domain, i.e.,

$$g(\mathbf{L})\mathbf{x} = g(\mathbf{U} \Lambda \mathbf{U}^H)\mathbf{x} = \mathbf{U} g(\Lambda) \mathbf{U}^H \mathbf{x} = \mathbf{U} g(\Lambda) \hat{\mathbf{x}}. \quad (1)$$

Spectral Clustering (SC) as low-pass filtering. SC is a well-known method for clustering nodes in the spectral domain. A simplified SC algorithm is described in the Appendix. Classical SC can be interpreted as low-pass filtering on a one-hot node signal $\delta_i \in \mathbb{R}^N$, where the i -th element of δ_i is 1, i.e. $\delta_i(i) = 1$, and 0 elsewhere, for each node i . Filtering δ_i in the graph Fourier domain can be expressed as

$$\mathbf{f}_i = g_{\lambda_L}(\Lambda) \mathbf{U}^H \delta_i, \quad (2)$$

where g_{λ_L} is the low-pass filter that filter out components whose frequencies are greater than λ_L as

$$g_{\lambda_L}(\lambda) = \begin{cases} 1 & \text{if } \lambda \leq \lambda_L; \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

As pointed out by Tremblay et al. (2016); Ramasamy and Madhoo (2015), the node distance $\|\mathbf{f}_i - \mathbf{f}_j\|^2$ used in SC can be approximated by filtering random node features with g_{λ_L} . Consider a random node feature matrix $\mathbf{R} = [\mathbf{r}_1 | \mathbf{r}_2 | \dots | \mathbf{r}_\eta] \in \mathbb{R}^{N \times \eta}$ consisting of η random features $\mathbf{r}_i \in \mathbb{R}^N$ i.i.d sampled from the normal distribution with zero mean and $1/\eta$ variance, let $\mathbf{H}_{\lambda_L} = \mathbf{U} g_{\lambda_L}(\Lambda) \mathbf{U}^H$, we can define

$$\tilde{\mathbf{f}}_i = (\mathbf{H}_{\lambda_L} \mathbf{R})^H \delta_i. \quad (4)$$

\mathbf{R} is a random Gaussian matrix of zero mean and \mathbf{U} is orthonormal. We apply the Johnson-Lindenstrauss lemma (see Appendix) to obtain the following error bounds.

Proposition 2.1 (Tremblay et al. (2016)). *Let $\epsilon, \beta > 0$ be given. If η is larger than:*

$$\eta_0 = \frac{4 + 2\beta}{\epsilon^2/2 - \epsilon^3/3} \log N, \quad (5)$$

then with probability at least $1 - N^{-\beta}$, we have: $\forall (i, j) \in [1, N]^2$,

$$(1 - \epsilon) \|\mathbf{f}_i - \mathbf{f}_j\|^2 \leq \|\tilde{\mathbf{f}}_i - \tilde{\mathbf{f}}_j\|^2 \leq (1 + \epsilon) \|\mathbf{f}_i - \mathbf{f}_j\|^2. \quad (6)$$

Hence, $\|\tilde{\mathbf{f}}_i - \tilde{\mathbf{f}}_j\|$ is a close estimation of the Euclidean distance $\|\mathbf{f}_i - \mathbf{f}_j\|$. Note that the approximation error bounds in Equation 6 also hold for any band-pass filter \bar{g}_{e_1, e_2} :

$$\bar{g}_{e_1, e_2}(\lambda) = \begin{cases} 1 & \text{if } e_1 < \lambda \leq e_2 \\ 0 & \text{otherwise.} \end{cases} \quad (7)$$

where $0 \leq e_1 \leq e_2 \leq 2$ since the spectrum of a normalized Laplacian matrix has the range $[0, 2]$. This fact is especially useful when the band pass \bar{g}_{e_1, e_2} can be expressed by a specific functional form such as polynomials, since $\bar{\mathbf{H}}_{e_1, e_2} = \mathbf{U} \bar{g}_{e_1, e_2}(\Lambda) \mathbf{U}^H$ can be computed without eigen-decomposition.

3 Adaptive Spectral Clustering

Now we propose an adaptive method for spectral clustering, which aligns the clustering structure with node labels by learning the underlying frequency patterns. This empowers us to restructure a graph to improve graph homophily while preserving the original graph structures as much as possible.

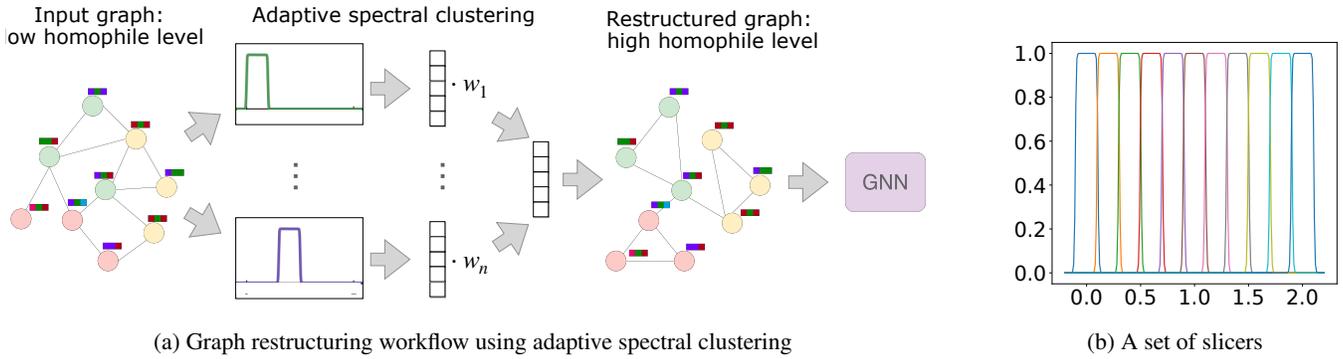


Figure 2: Adaptive spectral clustering using spectrum slicers. In 2a w_1, \dots, w_n are adaptive scalar parameters. In 2b, the spectrum range $[0, 2]$ is sliced into equal-width pseudo-eigenvalues by a set of slicers with $s = 20$.

Learning Eigenvector Coefficients

Let $\mathbf{f}_i^{\mathcal{Z}}$ be the representation of a node i obtained from an arbitrary set $\mathcal{Z} \subseteq [N]$ of eigenvectors in place of the ones with the leading eigenvalues. We cast the problem of finding eigenvectors to align with the known node labels into a minimization problem of computing the distance between the representations of nodes i and j when their labels are the same. Let $d(\cdot)$ be a distance metric between two node representations. Then, our objective is formalized as

$$\arg \min_{\mathcal{Z}} \sum_{i,j \in V_Y} \ell(d(\mathbf{f}_i^{\mathcal{Z}}, \mathbf{f}_j^{\mathcal{Z}}), \mathbb{1}(y_i, y_j)), \quad (8)$$

where V_Y is a collection of nodes whose labels are available, $\mathbb{1}$ is an indicator function, y_i is the label of node i , and ℓ is a loss function. The loss function penalizes the objective when two nodes of the same label are far from each other, as well as when two nodes of different labels are close. In this work, we use Euclidean distance and a contrastive loss function.

Solving the above objective requires iterating over all possible combinations of eigenvectors, which is infeasible in general. It also requires performing expensive eigendecomposition with $O(N^3)$ time complexity. In addition, classical SC does not consider node features when computing node distance. To address these challenges, we introduce two key ideas to generalize SC in the following.

Eigendecomposition-Free SC

As explained in Section 2, $\|\mathbf{f}_i - \mathbf{f}_j\|$ can be approximated by a filtering operation under the Johnson-Lindenstrauss lemma. The same holds true for the generalized case $\mathbf{f}_i^{\mathcal{Z}}$. However, the operation still requires expensive eigendecomposition as Equation 7 takes eigenvalues explicitly. To mitigate this issue, we propose to use a series of rectangular functions, each serving as a band-pass filter that ‘‘slices’’ the Laplacian spectrum into a finite set of equal-length and equal-magnitude ranges. Each filter takes the same form of Equation 7, but is relaxed on the continuous domain. This formulation comes with two major advantages. Firstly, rectangular functions can be efficiently approximated with polynomial rational functions, thus bypassing expensive eigendecomposition. Secondly, each g_{e_1, e_2} groups frequencies within the range $(e_1, e_2]$ to form a

‘‘coarsened’’ *pseudo-eigenvalue*. Because nearby eigenvalues capture similar structural information, the ‘‘coarsening’’ operation reduces the size of representations, while still largely preserving the representation power.

Spectrum slicers. We approximate the band-pass rectangular filters in Equation 7 using a rational function

$$\hat{g}_{s,a} = \frac{1}{s^{2m}} \left(\left(\frac{\lambda - a}{2 + \hat{\epsilon}} \right)^{2m} + \frac{1}{s^{2m}} \right)^{-1} \quad (9)$$

where $s \geq 2$ is a parameter that controls the width of the passing window on the spectrum, $a \in [0, 2]$ is a parameter that controls the horizontal center of the function, m is the approximation order. Figure 2b shows an example of these functions. With properly chosen s and a , the Laplacian spectrum can be evenly sliced into chunks of range $(\lambda_i, \lambda_{i+1})$. Each chunk is a *pseudo-eigenvalue* that umbrellas eigenvalues within the range. Substituting $g(\lambda)$ in Equation 1 with $\hat{g}_{s,a}$, the spectral filtering operation becomes

$$\mathbf{U} \hat{g}_{s,a}(\Lambda) \mathbf{U}^H \mathbf{x} = \frac{1}{s^{2m}} \left(\left(\frac{\mathbf{L} - a\mathbf{I}}{2 + \hat{\epsilon}} \right)^{2m} + \frac{\mathbf{I}}{s^{2m}} \right)^{-1}$$

where $\mathbf{U} \Lambda \mathbf{U}^H = \mathbf{L}$. An important property of Equation 10 is that the matrix inversion can be computed via truncated Neumann series (Wu et al. 2013). This can bring the computation cost of $O(N^3)$ down to $O(pN^2)$.

Lemma 1. For all $\hat{\epsilon} > \frac{2s^{2m}}{s^{2m}-1} - 2$, the inverse of $\mathbf{T} = \left(\frac{\mathbf{L} - a\mathbf{I}}{2 + \hat{\epsilon}} \right)^{2m} + \frac{\mathbf{I}}{s^{2m}}$ can be expressed by a Neumann series with guaranteed convergence (A proof is given in Appendix).

SC with Node Features

Traditionally, SC does not use node features. However, independent from graph structure, node features can provide extra information that is valuable for a clustering task (Bianchi, Grattarola, and Alippi 2020). We therefore incorporate node features into the SC filtering operation by concatenating it with the random signal before the filtering operation

$$\Gamma_{s,a} = \hat{g}_{s,a}(\mathbf{L})(\mathbf{R} \curvearrowright \mathbf{X}), \quad (10)$$

where \frown is a column-wise concatenation operation. $\mathbf{\Gamma}_{s,a}$ has the shape of $N \times (P + F)$ and is sometimes referred as "dictionary" (Thanou, Shuman, and Frossard 2014). When using a series of $g_{s,a}$, we have

$$\mathbf{\Gamma} = (\mathbf{\Gamma}_{s_1,a_1} \frown \mathbf{\Gamma}_{s_2,a_2} \frown \dots). \quad (11)$$

Let P' be the dimension of embeddings and Θ represent a weight matrix or a feed-forward neural network. The concatenated dictionary is then fed into a learnable function to obtain a node embedding matrix $\mathbf{H} \in \mathbb{R}^{N \times P'}$ as

$$\mathbf{H} = \Theta(\mathbf{\Gamma}). \quad (12)$$

Our objective in Equation 8 can then be realized as

$$\mathcal{L}(\Theta) = \sum_{\substack{i,j \in V_Y \\ k \in \mathcal{N}_Y(i) \\ y_i = y_j}} [\|\mathbf{H}_i - \mathbf{H}_j\|^2 - \|\mathbf{H}_i - \mathbf{H}_k\|^2 + \epsilon]_+ \quad (13)$$

where $\mathcal{N}_Y(k)$ is a set of negative samples whose labels are different from node i , i.e. $y_i \neq y_k$. The negative samples can be obtained by randomly sampling a fixed number of nodes with known labels. The intuition is if the labels of nodes i and j are the same, then the distance between the two nodes needs to be less than the distance between i and k , to minimize the loss. $[a]_+ = \max(a, 0)$ and ϵ is a scalar offset between distances of intra-class and inter-class pairs. By minimizing the objective w.r.t the weight matrix Θ , we can learn the weight of each band that aligns the best with the given labels. Ablation study and spectral expressive power of this method is discussed separately in Appendix.

Restructure Graphs to Maximize Homophily

After training, we obtain a distance matrix \mathbf{D}' where $D'_{ij} = \|\mathbf{H}_i - \mathbf{H}_j\|$. An intuitive way to reconstruct a graph is to start with a disconnected graph, and gradually add edges between nodes with the smallest distances, until the added edges do not increase homophily. We adopt the same approach as Klicpera, Weissenberger, and Günnemann (2019), to add edges between node pairs of the K smallest distance for better control of sparsity. Another possible way is to apply a threshold on \mathbf{D}' and entries above the threshold are kept as edges. Specifically, $\hat{\mathbf{A}}$ is the new adjacency matrix whose entries are defined as

$$\hat{A}_{ij} = \begin{cases} 1, & \text{if } (i, j) \in \text{topK}^{-1}(\mathbf{S}) \\ 0, & \text{otherwise,} \end{cases} \quad (14)$$

where topK^{-1} returns node pairs of the k smallest entries in \mathbf{D}' . A simplified workflow is illustrated in Figure 2a. A detailed algorithm can be found in Appendix.

Complexity Analysis

The most expensive operation of our method is the matrix inversion in Equation 10 which has the time complexity of $O(pN^2)$. A small $p \leq 4$ is sufficient because the Neumann series is a geometric sum so exponential acceleration tricks can be applied. Equation 10 is a close approximation to a rectangular function that well illustrates the spectrum slicing

idea. In practice, it can be replaced by other slicer functions that do not require matrix inversion, such as a quadratic function $1 - (s\mathbf{L} - a)^2$, to further reduce cost. It is also worth noting that the matrix inversion and multiplication only need to be computed once and can be pre-computed offline as suggested by (Rossi et al. 2020). The training step can be mini-batched easily. We randomly sample 8–64 negative samples per node so the cost of computing Equation 13 is low.

4 A New Homophily Measure

Several methods have been proposed to measure the homophily of a graph (Zhu et al. 2020; Lim et al. 2021). The two most used are edge homophily h_{edge} and node homophily h_{node} : the former uses the proportion of edges connecting nodes of the same label $h_{\text{edge}} = \frac{|(u,v) \in E: y_u = y_v|}{|E|}$, while the later uses the proportion of a node's direct neighbours of the same label $h_{\text{node}} = \frac{1}{N} \sum_{u \in V} \frac{|v \in \mathcal{N}_u: y_u = y_v|}{|\mathcal{N}_u|}$ where \mathcal{N}_u is the neighbour set of node u , and y is the node label.

As pointed out by Lim et al. (2021), both edge and node homophily suffer from sensitivity to both label number and label imbalance. For instance, a balanced graph of C labels would induce score of $\frac{1}{C}$ under both measures. In addition, both metrics fail to handle cases of imbalanced labels, resulting in undesirably high homophily scores when the majority nodes are of the same label.

To mitigate these issues, Lim et al. (2021) proposed a new metric h_{norm} that takes into account the label-wise node and edge proportion:

$$h_{\text{norm}} = \frac{1}{K-1} \sum_{k=0}^{K-1} \left[h_k - \frac{|Y_k|}{N} \right]_+, \quad (15)$$

$$(16)$$

$$h_k = \frac{\sum_{u \in Y_k} |v \in \mathcal{N}_u: y_u = y_v = k|}{\sum_{u \in Y_k} |\mathcal{N}_u|}, \quad (17)$$

K is the number of unique labels, Y_k is the node set of label k , and h_k is the label-wise homophily. Nevertheless, h_{norm} only captures relative edge proportions and ignores graph connectivity, resulting in high homophily scores for highly disconnected graphs. For example, Figure 3a has the same h_{norm} as 3b and 3c. The absence of edge density in the homophilic metric brings undesirable results in restructuring as the measurement always prefer a disconnected graph. Moreover, although h_{norm} is lower-bounded by 0, it does not explicitly define the meaning of $h_{\text{norm}} = 0$ but instead refers to such graphs as less-homophilic in general, resulting in further confusion when comparing less-homophilic graphs.

Given the limitations of existing metrics, we propose a density-aware homophily metric. For a graph of $K > 1$ labels, the following five propositions hold for our new metric:

1. A graph with a complete set of intra-class edges and zero inter-class edges has a score of 1. (Figure 3c)
2. A graph with a complete set of inter-class edges and zero intra-class edges has a score of 0. (Figure 3e)

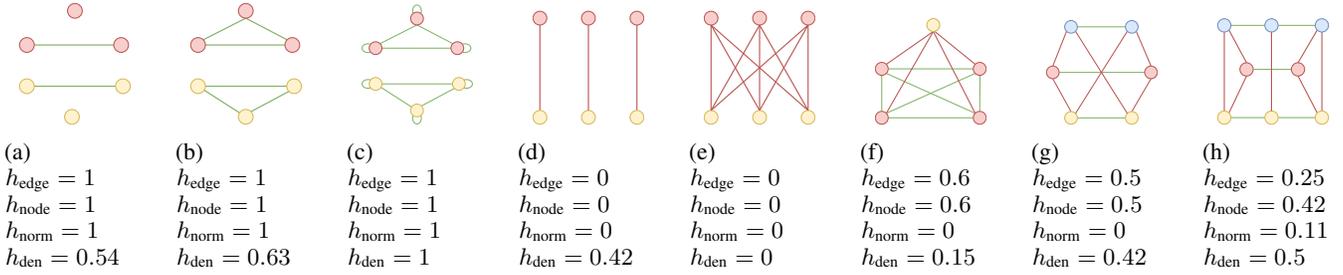


Figure 3: Examples of graphs with different label-topology relationships and comparison of different homophily measures. The node colour represents the node labels. The red/green edges connect nodes of different/same labels. Figure 3a - 3c shows homophilic graphs of different densities. h_{den} gives a higher score when a graph is dense, while the other metrics give the same scores. Figure 3d and 3e are two graphs that only consist of inter-class edges, but are of different densities. Figure 3f is a label-imbalanced graph. Figure 3g and Figure 3h are two regular graphs, where Figure 3g has an intra-class/inter-class edge ratio of 0.5, Figure 3h is an example of an Erdos-Reyi graph sampled with uniform edge probability.

3. An Erdos-Renyi random graph $G(n, p)$ has the score of ≈ 0.5 (Figure 3h).
4. A totally disconnected graph and a complete graph have the same score of 0.5.
5. For graphs with the same intra-class and inter-class edge ratios, the denser graph has a relatively higher score. (Figure 3a, 3b and 3c, Figure 3d and 3e, Figure 3g and 3h)

Propositions 1 and 2 define the limits of homophily and heterophily. Propositions 3 and 4 define neutral graphs which are neither homophilic nor heterophilic. Proposition 3 states that a uniformly random graph is neutral thus has a homophily score of 0.5. Proposition 5 considers edge density for graphs with the same inter- to intra-class edge ratio. The metric is defined as below.

$$\hat{h}_{\text{den}} = \min\{d_k - \bar{d}_k\}_{k=0}^{K-1} \quad (18)$$

where d_k is the edge density of the subgraph formed by nodes of label k , i.e. the intra-class edge density of k . \bar{d}_k is the maximum inter-class edge density of label k

$$d_k = \frac{|(u, v) \in E : k_u = k_v = k|}{|Y_k||Y_k|}, \quad (19)$$

$$\bar{d}_k = \max\{d_{k_j} : j = 0, \dots, K-1; j \neq k\}, \quad (20)$$

where d_{k_j} is the inter-class edge density of label j and k , i.e. edge density of the subgraph formed by nodes of label k or j .

$$d_{k_j} = \frac{|(u, v) \in E : k_u = k, k_v = j|}{|Y_k||Y_j|}. \quad (21)$$

We scale the metric to the range (0, 1) using

$$h_{\text{den}} = \frac{1 + \hat{h}_{\text{den}}}{2}. \quad (22)$$

For Proposition 3 we have the below lemma.

Lemma 2. $\forall K > 1, \mathbb{E}[h_{\text{den}}] = 0.5$ for the Erdos-Renyi random graph $G(n, p)$. (A proof is given in Appendix.)

Figure 3 shows some example graphs with four different metrics. Comparing with h_{node} and h_{edge} , h_{den} is not sensitive to the number of labels and label imbalance. Comparing with h_{norm} , h_{den} is able to detect neutral graphs. h_{den} gives scores in the range (0, 0.5) for graphs of low-homophily, allowing direct comparison between them.

5 Empirical Results

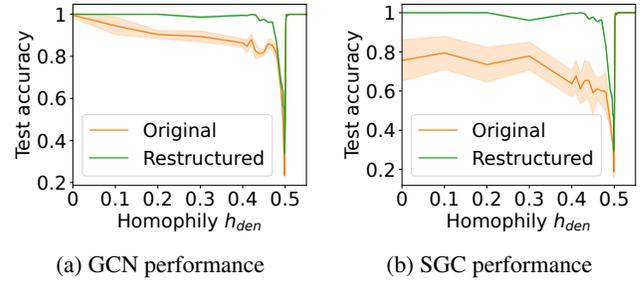


Figure 4: Performance on synthetic datasets.

Datasets and models. We compare six classical GNNs: GCN (Kipf and Welling 2017), SGC (Wu et al. 2019), ChevNet (Defferrard, Bresson, and Vandergheynst 2016), ARMANet (Bianchi et al. 2021), GAT (Velickovic et al. 2018), and APPNP (Klicpera, Bojchevski, and Günnemann 2019), on performance before and after restructuring. We also report performance using the restructuring methods GDC (Klicpera, Weißenberger, and Günnemann 2019). Five recent GNNs that target heterophilic graphs are also listed as baselines: GPRGNN (Chien et al. 2021), H₂GCN (Zhu et al. 2020), Geom-GCN (Pei et al. 2020), BernNet (He et al. 2021) and PPGNN (Lingam et al. 2021). We evaluate on four real-world graphs: TEXAS, CORNELL, CHAMELEON and SQUIRREL (Rozemberczki, Allen, and Sarkar 2021), as well as synthetic graphs of controlled homophily.

Experimental setup. Hyperparameters are tuned using grid search for all models on the unmodified and restructured graphs of each dataset. We record prediction accuracy on the test set averaged over 10 runs with different random initializations. We use the same split setting as Pei et al. (2020); Zhu et al. (2020). The results are averaged over all splits. We adopt early stopping and record the results from the epoch with highest validation accuracy. We report the averaged accuracy as well as the standard deviation. For the spectrum slicer in Equation 10, we use a set of 20 slicers

	CHAMELEON	SQUIRREL
GCN	59.8 ± 2.6	36.9 ± 1.3
GCN (GDC)	62.2 ± 1.2 (+2.4)	45.3 ± 1.3 (+8.4)
GCN (ours)	<u>66.9</u> ± 3.1 (+7.1)	<u>55.7</u> ± 2.4 (+18.8)
CHEV	66.0 ± 2.3	39.6 ± 3.0
CHEV(GDC)	63.0 ± 1.1 (-3.0)	48.2 ± 0.7 (+8.6)
CHEV(Ours)	<u>66.8</u> ± 1.8 (+0.8)	<u>55.0</u> ± 2.0 (+15.4)
ARMA	62.1 ± 3.6	47.8 ± 3.5
ARMA (GDC)	60.2 ± 0.6 (-1.9)	47.8 ± 0.8 (+0.0)
ARMA (ours)	<u>68.4</u> ± 2.3 (+6.3)	<u>55.6</u> ± 1.7 (+7.8)
GAT	54.7 ± 2.0	30.6 ± 2.1
GAT (GDC)	63.8 ± 1.2 (+9.1)	48.6 ± 2.1 (+18.0)
GAT (ours)	<u>66.5</u> ± 2.6 (+11.8)	<u>56.3</u> ± 2.2 (+25.7)
SGC	33.7 ± 3.5	46.9 ± 1.7
SGC (GDC)	60.6 ± 1.5 (+26.9)	51.4 ± 1.6 (+4.5)
SGC (ours)	<u>67.1</u> ± 2.9 (+33.4)	<u>52.3</u> ± 2.3 (+5.4)
APPNP	45.3 ± 1.6	31.0 ± 1.6
APPNP (GDC)	52.3 ± 1.4 (+7.0)	40.5 ± 0.8 (+9.5)
APPNP (ours)	<u>66.7</u> ± 2.7 (+21.4)	<u>55.9</u> ± 2.9 (+24.9)
GPRGNN	64.4 ± 1.6	41.9 ± 2.2
GPRGNN (GDC)	61.9 ± 1.7 (-2.5)	39.2 ± 1.5 (-1.7)
GPRGNN (ours)	<u>65.5</u> ± 2.2 (+1.1)	<u>47.1</u> ± 2.4 (+5.2)
Geom-GCN [†]	60.9	38.1
H ₂ GCN [*]	59.4 ± 2.0	37.9 ± 2.0
BernNet	62.0 ± 2.3 [∇]	37.9 ± 2.0
PPGNN	67.7 ± 2.3 [∇]	37.9 ± 2.0

Table 1: Node classification accuracy: CHAMELEON and SQUIRREL. Results marked with †, * and ∇ are obtained from Pei et al. (2020); Zhu et al. (2020); Lingam et al. (2021) respectively. Statistically significant results are underlined based on paired T-test of $p < 0.01$.

with $s = 40$ and $m = 4$ so that the spectrum is sliced into 20 even range of 0.1. In the restructure step, we add edges gradually and stop at the highest h_{den} on validation sets before the validation homophily starts to decrease. All experiments are run on a single NVIDIA RTX A6000 48GB GPU unless otherwise noted.

Node classification results. Node classification tasks predict labels of nodes based on graph structure and node features. We aim to improve the prediction accuracy of GNN models by restructuring edges via the adaptive SC method, particularly for heterophilic graphs. The evaluation results are shown in Table 1 and 2. On average, the performance of GNN models is improved by 25%. Training runtime for each dataset is reported in Appendix.

Results on synthetic graphs We run GCN and SGC on the synthetic dataset of controlled homophily range from 0 to 1. The model performance with homophily is plotted in Equation 4. As expected, higher homophily level corresponds to better performance for both GCN and SGC. All model reaches 100% accuracy where homophily is larger than 0.55. Performance reaches the lowest where homophily level is around 0.5. This is where intra-class edge has same density as inter-class edges, hence hard for GNN models to

	CORNELL	TEXAS
GCN	59.2 ± 3.2	64.1 ± 4.9
GCN (GDC)	57.6 ± 4.1 (-1.6)	57.8 ± 4.1 (-6.3)
GCN (ours)	<u>79.2</u> ± 6.3 (+20.0)	<u>78.4</u> ± 5.4 (+14.3)
CHEV	76.5 ± 9.4	79.7 ± 5.0
CHEV(GDC)	<u>81.1</u> ± 3.2 (+4.6)	79.2 ± 3.0 (-0.5)
CHEV(Ours)	80.8 ± 4.1 (+4.3)	<u>80.0</u> ± 4.8 (+0.3)
ARMA	74.9 ± 2.9	82.2 ± 5.1
ARMA (GDC)	78.4 ± 4.1 (+3.5)	78.4 ± 3.2 (-3.8)
ARMA (ours)	<u>81.1</u> ± 6.1 (+6.2)	<u>81.1</u> ± 4.2 (-1.1)
GAT	58.9 ± 3.3	60.0 ± 5.7
GAT (GDC)	58.9 ± 2.2 (+0.0)	77.1 ± 8.3 (+17.1)
GAT (ours)	<u>81.9</u> ± 5.4 (+23.0)	<u>79.8</u> ± 4.3 (+19.8)
SGC	58.1 ± 4.6	58.9 ± 6.1
SGC (GDC)	56.2 ± 3.8 (-1.9)	60.3 ± 6.3 (+1.4)
SGC (ours)	<u>73.5</u> ± 4.3 (+15.4)	<u>74.4</u> ± 6.0 (+15.5)
APPNP	70.3 ± 9.3	79.5 ± 4.6
APPNP (GDC)	77.8 ± 3.5 (+7.5)	76.2 ± 4.6 (-3.3)
APPNP (ours)	<u>81.6</u> ± 5.4 (+11.3)	<u>80.3</u> ± 4.8 (+0.8)
GPRGNN	79.5 ± 7.0	<u>84.6</u> ± 4.0
GPRGNN (GDC)	<u>82.4</u> ± 4.7 (+2.9)	80.8 ± 4.9 (-3.8)
GPRGNN (ours)	80.3 ± 6.3 (+0.8)	84.3 ± 5.1 (-0.3)
Geom-GCN [†]	60.8	67.6
H ₂ GCN [*]	82.2 ± 6.0	84.9 ± 6.8
BernNet	80.3 ± 5.4 [∇]	83.2 ± 6.5 [∇]
PPGNN	82.4 ± 4.3 [∇]	89.7 ± 4.9 [∇]

Table 2: Node classification accuracy: CORNELL and TEXAS.

capture useful structural patterns. Complete graphs, disconnected graphs and Erdos-erényi random graphs all fall into this category. When homophily level continues to decrease, performance starts to climb again. When homophily level reaches around 0, GNN models are able to perform as good as when homophily level is high. An explanation is, in such low-homophily cases, a node v aggregates features from nodes of other labels except its own, this aggregated information forms a useful pattern for predicting the label of v . Related observations are also reported by Luan et al. (2021); Ma et al. (2022). This shows traditional GNNs are actually able to capture the "opposite attract" pattern in extreme heterophilic graphs. In less-homophilic graphs where the homophily is less 0.5, our adaptive spectral clustering method is able to lift homophily to a very high level, leading to boosted performance. We also note that the performance variance on the rewired graph is much lower than these on the original graph. We detail the synthetic data generation process in Appendix.

Restructuring based on homophily. As homophily and performance are correlated, in the restructuring process, number of edges are chosen based on homophily level from validation set. As shown in Equation 5, we chose 48000 edges for CHAMELEON and 26000 edges for SQUIRREL, each corresponds to the first peak of homophily from the validation set. A limitation of the method is its reliance on a balanced data split, i.e. if the homophily on validation set does not represent the whole graph homophily, the restructured graph may not yield a good performance.

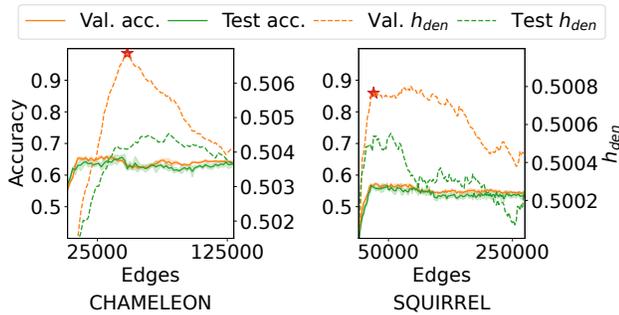


Figure 5: Homophily and accuracy of GCN on validation and test sets as per edges numbers. The optimal number of edges (red stars) are chosen based on h_{den} on validation set.

6 Related Work

GNNs for heterophilic graphs. Early GNNs assume homophily implicitly. Such an inductive bias results in a degenerated performance on less-homophilic graphs (Lim et al. 2021). Recently, homophily is also shown to be an effective measure of a graph’s robustness to both over-smoothing and adversarial attacks. Node representations converge to a stationary state of similar values (“over-smoothed”) as a GNN goes deeper. A graph of low homophily is also more prone to this issue as the stationary state is reached with fewer GNN layers (Yan et al. 2021). For the same reason, homophilic graphs are more resilient to a graph injection attack than their heterophilic counterparts. Some techniques defend against such attacks by improving or retaining homophily under graph injection (Zhu et al. 2021a; Chen et al. 2022). Pei et al. (2020) firstly draw attention to the limitation of GNN on less-homophilic graphs. Since then, various GNNs have been proposed to improve performance on these graphs. H_2GCN (Zhu et al. 2020) show that proper utilization of ego-embedding, higher-order neighbourhoods, and intermediate embeddings can improve results in heterophilic graphs. A recent scalable model, LINKX (Lim et al. 2021), shows separating feature and structural embedding improves performance. Kim and Oh (2021) study this topic specifically for graph attention and finds improvements when an attention mechanism is chosen according to homophily and average degrees. Chien et al. (2021) propose to use a generalized PageRank method that learns the weights of a polynomial filter and show that the model can adapt to both homophilic and heterophilic graphs. Similarly, Li, Kim, and Wang (2021) use learnable spectral filters for achieving an adaptive model on graphs of different homophilic levels. Zhu et al. (2021b) recently propose to incorporate a learnable compatibility matrix to handle heterophily of graphs. Ma et al. (2022) and Luan et al. (2021) found that when node neighbourhoods meet some condition, heterophily does not harm GNN performance. However, (Suresh et al. 2021) shows real-world graphs often exhibit mixing patterns of feature and neighbourhood proximity so such “good” heterophily is rare.

Adaptive spectral clustering. From the deep learning perspective, most previous studies about spectral clustering aim at clustering using learning approaches, or building relations

between a supervised learning model and spectral clustering. Law, Urtasun, and Zemel (2017) and Bach and Jordan (2003) reveal that minimizing the loss of node similarity matrix can be seen as learning the leading eigenvector representations used for spectral clustering. Bianchi, Grattarola, and Alippi (2020) train cluster assignment using a similarity matrix and further use the learned clusters in a pooling layer to for GNNs. Instead of directly optimizing a similarity matrix, Tian et al. (2014) adopt an unsupervised autoencoder design that uses node representation learned in hidden layers to perform K-means. Chowdhury and Needham (2021) show that Gromov-Wasserstein learning, an optimal transport-based clustering method, is related to a two-way spectral clustering.

Graph restructuring and rewiring. GDC (Klicpera, Weißenberger, and Günnemann 2019) is one of the first works propose to rewire edges in a graph. It uses diffusion kernels, such as heat kernel and personalized PageRank, to redirect messages passing beyond direct neighbours. Chamberlain et al. (2021) and Eliasof, Haber, and Treister (2021) extend the diffusion kernels to different classes of partial differential equations. Topping et al. (2022) studies the “over-squashing” issue on GNNs from a geometric perspective and alleviates the issue by rewiring graphs. In a slightly different setting where graph structures are not readily available, some try to construct graphs from scratch instead of modifying the existing edges. Fatemi, Asri, and Kazemi (2021) construct homophilic graphs via self-supervision on masked node features. Kalofolias (2016) learns a smooth graph by minimizing $tr(\mathbf{X}^T \mathbf{L} \mathbf{X})$. These works belong to the “learning graphs from data” family (Dong et al. 2019) which is relevant but different to our work because *a)* these methods infer graphs from data where no graph topology is readily available, while in our setting, the original graph is a key ingredient; *b)* as shown by Fatemi, Asri, and Kazemi (2021), without the initial graph input, the performance of these methods are not comparable to even a naive GCN; *c)* these methods are mostly used to solve graph generation problems instead of node- or link-level tasks.

7 Conclusion

We propose an approach to enhance GNN performance on less-homophilic graphs by restructuring the graph to maximize homophily. Our method is inspired and closely related to Spectral Clustering (SC). It extends SC beyond the leading eigenvalues and learns the frequencies that are best suited to cluster a graph. To achieve this, we use rectangular spectral filters expressed in the Neumann series to slice the graph spectrum into chunks that umbrellas small ranges of frequency. We also proposed a new homophily metric that is density-aware, and robust to label imbalance, hence a better homophily indicator for the purpose of graph restructuring¹. There are many promising extensions of this work, such as using it to guard against over-smoothing and adversarial attacks, by monitoring changes in homophily and adopting tactics to maintain it. We hereby leave these as future work.

¹The extended version is available at <https://arxiv.org/abs/2206.02386>. The code is available at https://github.com/seanli3/graph_restructure.

Acknowledgements

This work was partly supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No.2019-0-01906, Artificial Intelligence Graduate School Program(POSTECH)) and National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (NRF-2021R1C1C1011375). Dongwoo Kim is the corresponding author.

References

- Bach, F.; and Jordan, M. 2003. Learning spectral clustering. *Advances in neural information processing systems*, 16.
- Bianchi, F. M.; Grattarola, D.; and Alippi, C. 2020. Spectral clustering with graph neural networks for graph pooling. In *International Conference on Machine Learning*, 874–883. PMLR.
- Bianchi, F. M.; Grattarola, D.; Livi, L.; and Alippi, C. 2021. Graph neural networks with convolutional arma filters. *IEEE transactions on pattern analysis and machine intelligence*.
- Bo, D.; Wang, X.; Shi, C.; and Shen, H. 2021. Beyond Low-frequency Information in Graph Convolutional Networks. In *Thirty-Fifth AAAI Conference on Artificial Intelligence*.
- Chamberlain, B.; Rowbottom, J.; Gorinova, M. I.; Bronstein, M. M.; Webb, S.; and Rossi, E. 2021. GRAND: Graph Neural Diffusion. In *Proceedings of the 38th International Conference on Machine Learning, ICML*.
- Chen, Y.; Yang, H.; Zhang, Y.; Ma, K.; Liu, T.; Han, B.; and Cheng, J. 2022. Understanding and Improving Graph Injection Attack by Promoting Unnoticeability. *CoRR*, abs/2202.08057.
- Chien, E.; Peng, J.; Li, P.; and Milenkovic, O. 2021. Adaptive Universal Generalized PageRank Graph Neural Network. In *9th International Conference on Learning Representations, ICLR*.
- Chowdhury, S.; and Needham, T. 2021. Generalized Spectral Clustering via Gromov-Wasserstein Learning. In *The 24th International Conference on Artificial Intelligence and Statistics, AISTATS*.
- Defferrard, M.; Bresson, X.; and Vandergheynst, P. 2016. Convolutional Neural Networks on Graphs with Fast Localized Spectral Filtering. In *Advances in Neural Information Processing Systems (NeurIPS)*, 3837–3845.
- Dong, X.; Thanou, D.; Rabbat, M.; and Frossard, P. 2019. Learning Graphs from Data: A Signal Representation Perspective. *IEEE Signal Process. Mag.*, 36(3): 44–63.
- Eliasof, M.; Haber, E.; and Treister, E. 2021. PDE-GCN: Novel architectures for graph neural networks motivated by partial differential equations. *Advances in Neural Information Processing Systems*, 34.
- Fatemi, B.; Asri, L. E.; and Kazemi, S. M. 2021. SLAPS: Self-Supervision Improves Structure Learning for Graph Neural Networks. In *Advances in Neural Information Processing Systems 34: NeurIPS 2021*.
- He, M.; Wei, Z.; Huang, Z.; and Xu, H. 2021. BernNet: Learning Arbitrary Graph Spectral Filters via Bernstein Approximation. In Ranzato, M.; Beygelzimer, A.; Dauphin, Y. N.; Liang, P.; and Vaughan, J. W., eds., *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, 14239–14251.
- Kalofolias, V. 2016. How to Learn a Graph from Smooth Signals. *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics, AISTATS*.
- Kim, D.; and Oh, A. 2021. How to Find Your Friendly Neighborhood: Graph Attention Design with Self-Supervision. In *9th International Conference on Learning Representations, ICLR*.
- Kipf, T. N.; and Welling, M. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *Proceedings of the 5th International Conference on Learning Representations (ICLR)*.
- Klicpera, J.; Bojchevski, A.; and Günnemann, S. 2019. Predict then Propagate: Graph Neural Networks meet Personalized PageRank. In *Proceedings of the 7th International Conference on Learning Representations (ICLR)*.
- Klicpera, J.; Weißenberger, S.; and Günnemann, S. 2019. Diffusion Improves Graph Learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 13333–13345.
- Law, M. T.; Urtasun, R.; and Zemel, R. S. 2017. Deep Spectral Clustering Learning. In *Proceedings of the 34th International Conference on Machine Learning, ICML*. PMLR.
- Li, S.; Kim, D.; and Wang, Q. 2021. Beyond Low-Pass Filters: Adaptive Feature Propagation on Graphs. In *Machine Learning and Knowledge Discovery in Databases. Research Track - European Conference, ECML PKDD*.
- Lim, D.; Hohne, F.; Li, X.; Huang, S. L.; Gupta, V.; Bhalerao, O.; and Lim, S. 2021. Large Scale Learning on Non-Homophilous Graphs: New Benchmarks and Strong Simple Methods. In *Advances in Neural Information Processing Systems 34: NeurIPS*.
- Lingam, V.; Ekbote, C.; Sharma, M.; Ragesh, R.; Iyer, A.; and Sellamanickam, S. 2021. A Piece-wise Polynomial Filtering Approach for Graph Neural Networks. In *ICLR Workshop 2022*.
- Luan, S.; Hua, C.; Lu, Q.; Zhu, J.; Zhao, M.; Zhang, S.; Chang, X.; and Precup, D. 2021. Is Heterophily A Real Nightmare For Graph Neural Networks To Do Node Classification? *CoRR*, abs/2109.05641.
- Ma, Y.; Liu, X.; Shah, N.; and Tang, J. 2022. Is Homophily a Necessity for Graph Neural Networks? In *Proceedings of the 10th International Conference on Learning Representations (ICLR)*.
- Miller, B. A.; Beard, M. S.; and Bliss, N. T. 2011. Matched filtering for subgraph detection in dynamic networks. In *2011 IEEE Statistical Signal Processing Workshop (SSP)*, 509–512.
- Pei, H.; Wei, B.; Chang, K. C.; Lei, Y.; and Yang, B. 2020. Geom-GCN: Geometric Graph Convolutional Networks. In *8th International Conference on Learning Representations, ICLR*.

- Ramasamy, D.; and Madhoo, U. 2015. Compressive spectral embedding: sidestepping the SVD. In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems*.
- Ribeiro, L. F.; Saverese, P. H.; and Figueiredo, D. R. 2017. Struc2vec: Learning node representations from structural identity. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- Rossi, E.; Frasca, F.; Chamberlain, B.; Eynard, D.; Bronstein, M. M.; and Monti, F. 2020. SIGN: Scalable Inception Graph Neural Networks. In *Proceedings of the 37th International Conference on Machine Learning, ICML*.
- Rozemberczki, B.; Allen, C.; and Sarkar, R. 2021. Multi-Scale attributed node embedding. *J. Complex Networks*, 9(2).
- Schaub, M. T.; and Segarra, S. 2018. Flow Smoothing and Denoising: Graph Signal Processing in the Edge-Space. In *2018 IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, 735–739.
- Shuman, D. I.; Narang, S. K.; Frossard, P.; Ortega, A.; and Vandergheynst, P. 2013. The Emerging Field of Signal Processing on Graphs: Extending High-Dimensional Data Analysis to Networks and Other Irregular Domains. *IEEE Signal Process. Mag.*, 30(3): 83–98.
- 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 Zhu, F.; Ooi, B. C.; and Miao, C., eds., *KDD '21: The 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Virtual Event, 2021*.
- Thanou, D.; Shuman, D. I.; and Frossard, P. 2014. Learning parametric dictionaries for signals on graphs. *IEEE Trans. Signal Process.*, 62(15): 3849–3862.
- Tian, F.; Gao, B.; Cui, Q.; Chen, E.; and Liu, T. Y. 2014. Learning deep representations for graph clustering. *Proceedings of the National Conference on Artificial Intelligence*, 2: 1293–1299.
- Topping, J.; Giovanni, F. D.; Chamberlain, B. P.; Dong, X.; and Bronstein, M. M. 2022. Understanding over-squashing and bottlenecks on graphs via curvature. In *Proceedings of the 10th International Conference on Learning Representations (ICLR)*.
- Tremblay, N.; Puy, G.; Borgnat, P.; Gribonval, R.; and Vandergheynst, P. 2016. Accelerated spectral clustering using graph filtering of random signals. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP*.
- 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 (ICLR)*.
- Wai, H.; Segarra, S.; Ozdaglar, A. E.; Scaglione, A.; and Jadbabaie, A. 2018. Community Detection from Low-Rank Excitations of a Graph Filter. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 4044–4048.
- Wu, F.; Jr., A. H. S.; Zhang, T.; Fifty, C.; Yu, T.; and Weinberger, K. Q. 2019. Simplifying Graph Convolutional Networks. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, volume 97, 6861–6871.
- Wu, M.; Yin, B.; Vosoughi, A.; Studer, C.; Cavallaro, J. R.; and Dick, C. 2013. Approximate matrix inversion for high-throughput data detection in the large-scale MIMO uplink. In *2013 IEEE International Symposium on Circuits and Systems (ISCAS)*, 2155–2158.
- Yan, Y.; Hashemi, M.; Swersky, K.; Yang, Y.; and Koutra, D. 2021. Two sides of the same coin: Heterophily and over-smoothing in graph convolutional neural networks. *arXiv preprint arXiv:2102.06462*.
- Ying, R.; Bourgeois, D.; You, J.; Zitnik, M.; and Leskovec, J. 2019. GNNExplainer: Generating explanations for graph neural networks. *Adv. Neural Inf. Process. Syst.*, 32(July 2017).
- Zeng, H.; Zhou, H.; Srivastava, A.; Kannan, R.; and Prasanna, V. K. 2020. GraphSAINT: Graph Sampling Based Inductive Learning Method. In *8th International Conference on Learning Representations, ICLR*.
- Zhu, J.; Jin, J.; Schaub, M. T.; and Koutra, D. 2021a. Improving Robustness of Graph Neural Networks with Heterophily-Inspired Designs. *CoRR*, abs/2106.07767.
- Zhu, J.; Rossi, R. A.; Rao, A. B.; Mai, T.; Lipka, N.; Ahmed, N. K.; and Koutra, D. 2021b. Graph Neural Networks with Heterophily. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- 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 *Advances in Neural Information Processing Systems 33: NeurIPS*.