

Task-driven Heterophilic Graph Structure Learning

Ayushman Raghuvanshi¹, Gonzalo Mateos², and Sundeep Prabhakar Chepuri¹

¹ Indian Institute of Science, Bangalore, India. ² University of Rochester, Rochester, United States

Abstract—Graph neural networks (GNNs) often struggle to learn discriminative node representations for heterophilic graphs, where connected nodes tend to have dissimilar labels and feature similarity provides weak structural cues. We propose *frequency-guided graph structure learning* (FgGSL), an end-to-end graph inference framework that jointly learns homophilic and heterophilic graph structures along with a spectral encoder. FgGSL employs a learnable, symmetric, feature-driven masking function to infer said complementary graphs, which are processed using pre-designed low- and high-pass graph filter banks. A label-based structural loss explicitly promotes the recovery of homophilic and heterophilic edges, enabling task-driven graph structure learning. We derive stability bounds for the structural loss and establish robustness guarantees for the filter banks under graph perturbations. Experiments on six heterophilic benchmarks demonstrate that FgGSL consistently outperforms state-of-the-art GNNs and graph rewiring methods, highlighting the benefits of combining frequency information with supervised topology inference.

Index Terms—Graph learning, Graph filtering, Heterophilic graphs, Node classification, Topology inference.

I. INTRODUCTION

Graphs provide an effective framework for modeling relational patterns in complex domains, such as social networks, biochemical systems, and financial markets, where the connectivity among entities conveys valuable information about the entities themselves. Recent advances in graph representation learning (GRL) have enabled the integration of this structural information into parameterized predictive models, significantly improving performance on tasks such as node classification, link prediction, and node regression [1].

Most GRL methods, including message passing neural networks (MPNNs) and graph neural networks (GNNs), rely on the assumption of *homophily*, meaning that connected nodes tend to be semantically similar. This assumption aligns with aggregation-based and low-pass filtering approaches that promote smooth node embeddings. However, many real-world networks, such as protein–protein interaction and online discussion networks, exhibit *heterophily*, where connected nodes often have dissimilar labels [2]. In such cases, conventional GNNs designed under the homophilic assumption tend to perform poorly due to feature mixing between unrelated nodes [3].

Research on representation learning for heterophilic graphs aims to overcome the limitations imposed by the homophily assumptions underlying standard GNNs. Approaches include redefining architectures, aggregation mechanisms, and customizing spectral filters to capture broader frequency components of graph signals [2], [4]–[6]. These methods primarily enhance the GNN’s ability to exploit graph structure and handle heterophily more effectively.

Alternatively, *graph rewiring* approaches aim to modify the graph structure to make it compatible with legacy GNNs. These methods add homophilic edges, prune heterophilic edges, or use score-based strategies to yield a more suitable topology [7], [8]. However, their effectiveness is limited because most rewiring strategies rely primarily on node features to guide structural updates. In heterophilic settings, feature similarity is a poor proxy for label similarity, resulting in the introduction of edges that are not semantically meaningful, even after rewiring. As a result, many rewired graphs perform only marginally




Fig. 1: Distribution of cosine similarities of node features \mathbf{x} for node pairs in the Texas (left) and Cornell (right) heterophilic datasets. Homophilic pairs (same labels) and heterophilic pairs (different labels) both exhibit overlapping distributions, implying that feature similarity alone is insufficient to recover graph structures.

better than, or even comparably to, graph-agnostic models such as multilayer perceptrons (MLP). This first challenge highlights the need for a principled, task-driven framework that can leverage label information to infer graph structure in heterophilic settings.

The goal of graph inference is to recover an underlying graph for given nodal features and subsequently apply GRL methods [9]. Smoothness-based graph inference approaches attempt to identify a graph that promotes feature smoothness while enforcing structural properties such as edge sparsity. However, these techniques also face challenges on heterophilic datasets, where node feature similarity does not correlate well with label similarity. In other words, same-label and different-label node pairs present highly overlapping feature-similarity distributions, as shown in Fig. 1, underscoring the lack of separability between intra-class and inter-class node pairs when relying solely on node features.

Contributions. To address the aforementioned twofold challenge, we propose the *frequency-guided graph structure learning* (FgGSL) framework that jointly learns task-aware homophilic and heterophilic graph structures in a supervised setting. FgGSL learns two distinct masking functions that generate homophilic and heterophilic graphs from a fully connected graph, and then applies distinct graph filters designed to capture complementary regions of the graph spectral domain. The homophilic graph leverages pre-designed low-pass filters to smooth out the features of connected nodes, while applying node label-based smoothness regularization to promote homophilic graph recovery. On the other hand, the heterophilic graph utilizes pre-designed high-pass filters to emphasize differences in features between connected nodes, and a node label-based anti-smoothness term that guides the recovery of heterophilic graphs. By integrating the filtered signals from these two complementary graphs, FgGSL captures the rich spectral characteristics of the node features in these complex datasets.

We further establish theoretical bounds on the deviation of the filtered output for graphs close to the optimal graph structure, following the approach in [10]. Hence, the model remains robust even if the recovered graph is sub-optimal, as long as it is close to the optimal structure. Since testing node labels are unavailable during training, we use predicted labels and guarantee that the structural loss remains close to the actual label-based counterpart, as long as the prediction error is bounded. This enables FgGSL to maintain its

This work is supported in part by the Kotak IISc AI-ML Center (KIAC).




Fig. 2: Illustration of a homophilic graph (left), and a heterophilic graph (right). The color of the nodes indicates their labels.

performance even without the exact labels of the testing node.

To evaluate the effectiveness of the proposed framework, we perform node classification experiments on heterophilic benchmarks [2]. The results demonstrate that FgGSL consistently outperforms existing GNNs designed for heterophilic settings as well as graph rewiring approaches. Moreover, the learned embedding exhibits a markedly different similarity distribution among the semantically similar and dissimilar node pairs. In summary, our work advances task-driven graph structure learning under heterophily by addressing the limitations of feature-only rewiring and (fixed-graph) GNNs via the following key contributions:

- We introduce a frequency-guided, task-driven graph inference framework that jointly learns homophilic and heterophilic graph structures tailored to downstream prediction tasks;
- We propose label-aware structural loss functions that explicitly encourage homophily and heterophily, enabling supervised as well as interpretable graph inference; and
- We develop an interpretable spectral architecture that combines predefined low-pass and high-pass filter banks to capture complementary frequency components in node representations.

We empirically validate that, by jointly modeling low- and high-frequency modes through complementary spectral filters, FgGSL effectively captures informative relationships that conventional aggregation-based GNNs fail to exploit.

II. BACKGROUND

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be an undirected graph, where \mathcal{V} is the node set with $|\mathcal{V}| = N$ and \mathcal{E} represents the set of edges. Each node $i \in \mathcal{V}$ is associated with a feature vector $\mathbf{x}_i \in \mathbb{R}^F$, and the collection of all node features is represented by the matrix $\mathbf{X} \in \mathbb{R}^{N \times F}$. One-hot encoded node labels \mathbf{y}_v for $v \in \mathcal{V}$ are stored in $\mathbf{Y} \in \{0, 1\}^{N \times C}$, where C is the number of classes. The adjacency matrix $\mathbf{A} \in \{0, 1\}^{N \times N}$ encodes the graph connectivity, where $A_{ij} = 1$ if $(i, j) \in \mathcal{E}$ and $A_{ij} = 0$, otherwise. The degree matrix is $\mathbf{D} = \text{diag}(\sum_j A_{ij})$, and the normalized Laplacian is given by $\mathbf{L} = \mathbf{I} - \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}$.

Heterophilic graphs. Graphs in which adjacent nodes are often semantically different, meaning they have more inter-class edges than intra-class edges, are known as *heterophilic graphs* (see Fig. 2). One way to quantify heterophily is via the *heterophilic edge ratio*, defined as the ratio of heterophilic edges to the total number of edges in the graph, i.e., $R_{\text{het}}(\mathcal{G}) = |\{(u, v) \in \mathcal{E} : \mathbf{y}_u \neq \mathbf{y}_v\}| / |\mathcal{E}|$.

Graph Fourier transform (GFT). The GFT is an orthonormal transform that decomposes a graph signal into different “frequency” modes of graph-dependent variation. Low frequencies correspond to smooth signals (connected nodes have similar values), while high frequencies imply signals exhibit high variability across neighbors. We leverage this spectral interpretation of graph signals to explicitly extract low-frequency (smooth) signals on the learned homophilic graph and high-frequency signals on the learned heterophilic graph.

For a graph \mathcal{G} whose Laplacian has eigendecomposition $\mathbf{L} = \mathbf{U} \mathbf{\Lambda} \mathbf{U}^\top$, the GFT of a graph signal \mathbf{X} is defined as $\tilde{\mathbf{X}} = \mathbf{U}^\top \mathbf{X}$, where $\tilde{\mathbf{x}}_i^\top = \mathbf{u}_i^\top \mathbf{X}$ denotes the spectral component of \mathbf{X} associated with eigenvalue (i.e., frequency) λ_i . Each eigenpair $(\lambda_i, \mathbf{u}_i)$ represents a frequency mode of the graph with smaller eigenvalues corresponding to smooth, low-frequency signal components.

A *graph filter* is a function of \mathbf{L} that can impose desired spectral characteristics onto a graph signal. It is specified in terms of a spectral kernel $h : \mathbb{R} \mapsto \mathbb{R}$, which modulates the frequency content of a signal. Filtering a graph signal \mathbf{X} with kernel $h(\cdot)$ boils down to

$$\tilde{\mathbf{X}} = h(\mathbf{L}) \mathbf{X} = \mathbf{U} h(\mathbf{\Lambda}) \mathbf{U}^\top \mathbf{X} = \mathbf{U} h(\mathbf{\Lambda}) \tilde{\mathbf{X}},$$

where $h(\mathbf{\Lambda}) = \text{diag}(h(\lambda_1), \dots, h(\lambda_N))$ is a diagonal matrix containing the spectral response of the filter. Here, $h(\lambda_i)$ determines the degree of amplification or attenuation applied to the spectral component corresponding to frequency λ_i .

Graph structure inference. The problem of graph inference can be separated into two stages: (i) unsupervised inference of a graph structure, often guided by smoothness assumptions and regularizers; and (ii) supervised training of a GNN on the inferred graph.

Formally, given a graph signal \mathbf{X} over nodes \mathcal{V} along with labels \mathbf{Y} , the goal in (i) is to learn an optimal graph Laplacian \mathbf{L}^* given a structural loss $\ell_{st} : \mathbb{R}^{N \times N} \times \mathbb{R}^{N \times f} \mapsto \mathbb{R}$ as

$$\mathbf{L}^* = \arg \min_{\mathbf{L} \in \mathcal{L}} \ell_{st}(\mathbf{L}, \mathbf{X}), \quad (1)$$

where \mathcal{L} is the set of valid Laplacians. Typically, for $\ell_{st}(\mathbf{L}, \mathbf{X})$, we use the quadratic smoothness $\text{Tr}(\mathbf{X}^\top \mathbf{L} \mathbf{X})$ for learning homophilic graphs or anti-smoothness $-\text{Tr}(\mathbf{X}^\top \mathbf{L} \mathbf{X})$ for learning heterophilic graphs. This is followed by (ii), namely,

$$\phi^* = \arg \min_{\phi} \|\mathbf{Y} - f_{\phi}(\mathbf{L}^*, \mathbf{X})\|, \quad (2)$$

where f_{ϕ} is a GNN with the learnable parameter ϕ .

Solving this problem directly is challenging for heterophilic graph signals because score functions based solely on node feature similarity, such as anti-smoothness is not useful as shown in Fig. 1, where recall that the feature-similarity distribution of same-class node pairs overlaps with that of different-class pairs. Consequently, the learned Laplacian \mathbf{L}^* constructed solely from node feature similarity does not necessarily exhibit heterophily.

III. PROPOSED MODEL: FREQUENCY-GUIDED GRAPH STRUCTURE LEARNING (FGGSL)

In this section, we reformulate (1) and (2) as a joint graph learning and encoder training problem. We introduce two label-based structural loss functions that incorporate label similarity into the structural score. We then develop a parameterized method to explore the graph search space, define filter banks to generate node representations, and introduce a label-based structural loss that guides the learning of strong heterophilic and homophilic graphs. We judiciously combine these ingredients to arrive at the novel FgGSL framework.

Problem formulation. Consider a heterophilic graph \mathcal{G} with node features \mathbf{X} and one-hot encoded class labels \mathbf{Y} . Our goal is to jointly learn an homophilic graph $\hat{\mathcal{G}}_{\text{Ho}}$, a heterophilic graph $\hat{\mathcal{G}}_{\text{Ht}}$, and the parameters $\hat{\phi}$ of an encoder f_{ϕ} by solving the optimization problem:

$$\{\hat{\mathcal{G}}_{\text{Ho}}, \hat{\mathcal{G}}_{\text{Ht}}, \hat{\phi}\} = \arg \min_{\mathcal{G}_{\text{Ho}}, \mathcal{G}_{\text{Ht}}, \phi} \ell(\mathcal{G}_{\text{Ho}}, \mathcal{G}_{\text{Ht}}, \mathbf{X}, f_{\phi}, \mathbf{Y}). \quad (3)$$

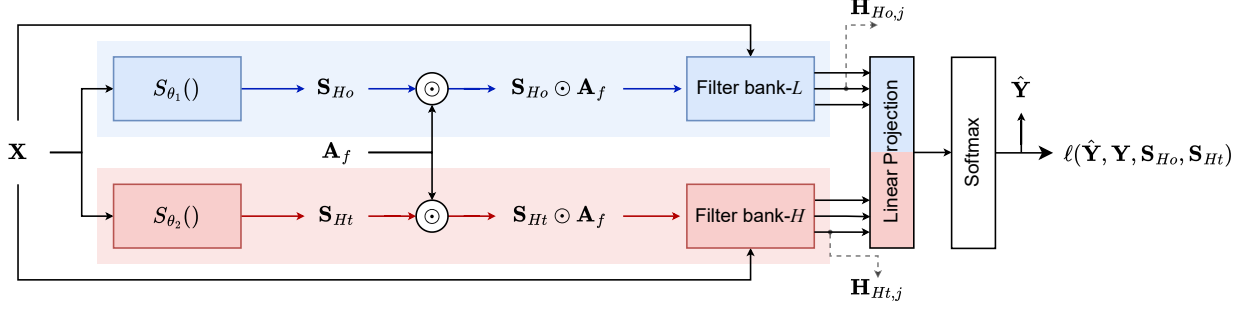


Fig. 3: **FgGSL framework**: Node features \mathbf{X} generate two parameterized graph structures by masking a fully connected adjacency matrix \mathbf{A}_f with learnable functions S_{θ_1} and S_{θ_2} . The homophilic graph $\mathbf{S}_{Ho} \odot \mathbf{A}_f$ uses a low-pass filter bank, while the heterophilic graph $\mathbf{S}_{Ht} \odot \mathbf{A}_f$ employs a high-pass filter bank. Filter outputs $\mathbf{H}_{Ho,j}$ and $\mathbf{H}_{Ht,j}$ at different scales are concatenated and passed through a linear layer with softmax to produce class probabilities $\hat{\mathbf{Y}}$.

Here, ℓ is the joint graph structure and encoder training loss function that takes the form

$$\begin{aligned} \ell &= \ell_{\text{CE}}(f_\phi(\mathcal{G}_{Ho}, \mathcal{G}_{Ht}, \mathbf{X}), \mathbf{Y}) + \ell_{\text{SL}}(\mathcal{G}_{Ho}, \mathcal{G}_{Ht}, \mathbf{Y}), \\ \ell_{\text{SL}} &= \alpha \ell_{Ho}(\mathcal{G}_{Ho}, \mathbf{Y}) + \beta \ell_{Ht}(\mathcal{G}_{Ht}, \mathbf{Y}), \end{aligned}$$

where $\alpha, \beta > 0$ are hyper-parameters. The cross-entropy loss ℓ_{CE} is the task-driven loss used to train the encoder; alternatively, other losses may be used depending on the downstream task. The second term ℓ_{SL} is the structural loss that includes both a homophilic loss ℓ_{Ho} , promoting label smoothness over \mathcal{G}_{Ho} , and a heterophilic loss ℓ_{Ht} , promoting label dissimilarity over \mathcal{G}_{Ht} ; see also (7)-(8) for details. The FgGSL framework, illustrated in Fig. 3, minimizes the loss by first parametrizing the graph structure via learnable masking, then applying spectral filtering using filter banks. The outputs of the filters are concatenated and linearly combined to produce the logits for node classification

Parameterized graph structure. To make the graph structure learnable, we express \mathcal{G}_{Ho} and \mathcal{G}_{Ht} as masked versions of a fully connected graph $\mathcal{G}_f = (\mathcal{V}, \mathcal{E}_f)$ with adjacency \mathbf{A}_f . This parameterization enables the model to adjust edge weights and effectively infer the optimal connectivity pattern based on node features. Specifically, we define a learnable masking function $S_\theta : \mathbb{R}^F \times \mathbb{R}^F \mapsto [0, 1]$ that assigns a weight $w_{ij} \in [0, 1]$ to each edge (i, j) in \mathcal{E}_f . This masking function is defined as the inner product between transformed node features. Therefore, for a parameterized nonlinear map $\Phi_\theta : \mathbb{R}^F \mapsto \mathbb{R}^D$, the mask is computed as:

$$S_\theta(\mathbf{x}_u, \mathbf{x}_v) = \sigma(\Phi_\theta^\top(\mathbf{x}_u), \Phi_\theta(\mathbf{x}_v)) = w_{uv}, \quad (u, v) \in \mathcal{E}_f, \quad (4)$$

where $\sigma(\cdot)$ denotes the sigmoid activation. Thus, for a given graph \mathcal{G}_f , we obtain a masked homophilic graph $\hat{\mathcal{G}}_{Ho} = (\mathcal{V}, \mathcal{E}_f, S_{\theta_1})$ and a heterophilic graph $\hat{\mathcal{G}}_{Ht} = (\mathcal{V}, \mathcal{E}_f, S_{\theta_2})$, with adjacencies $\mathbf{S}_{Ho} \odot \mathbf{A}_f$ and $\mathbf{S}_{Ht} \odot \mathbf{A}_f$, respectively. Therefore, the search space for the graph is parameterized by the learnable weights θ in the nonlinear map Φ_θ , which govern the edge masking process. In particular, we use parameters θ_1 for the homophilic graph and θ_2 for the heterophilic graph. Intuitively, edges with weights close to zero contribute minimally to the learned structure. When \mathcal{G}_f is fully connected, this formulation allows the model to infer any possible graph as a masked variant of the initial one. Alternatively, when \mathcal{G}_f corresponds to some graph available in the dataset, the masking operation effectively performs edge pruning. In practice, \mathcal{G}_f may be either fully connected or the given graph.

Filter banks. The function f_ϕ , optimized in (3), represents the classifier that produces class probabilities $\hat{\mathbf{Y}} \in [0, 1]^{N \times C}$ from

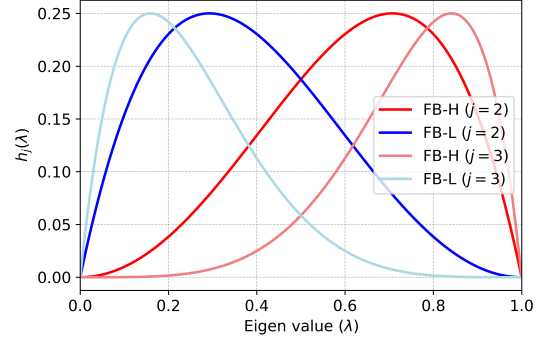


Fig. 4: Filter-bank frequency responses: Filters from filter bank-L preserve information from the lower end of the spectrum (blue), while filters from filter bank-H preserve information from the higher end of the spectrum (red).

node features \mathbf{X} and the learned graphs $\{\hat{\mathcal{G}}_{Ho}, \hat{\mathcal{G}}_{Ht}\}$. In FgGSL, the classifier employs two sets of graph filters: low-pass filters applied to the homophilic graph \mathcal{G}_{Ho} and high-pass filters applied to the heterophilic graph \mathcal{G}_{Ht} . The low-pass filters extract smooth, shared information among similar nodes, while the high-pass filters highlight discriminative or contrasting information among dissimilar nodes.

Specifically, we use predefined low-pass filter banks $\{h_L^{(j)}\}_{j=2}^J$ on \mathcal{G}_{Ho} and high-pass filter banks $\{h_H^{(j)}\}_{j=2}^J$ on \mathcal{G}_{Ht} . For a given maximum scale J , the filter kernels are defined as

$$h_L^{(j)}(\lambda) = (0.5\lambda)^{2^{j-1}} - (0.5)^{2^j}, \quad (5)$$

$$h_H^{(j)}(\lambda) = (1 - 0.5\lambda)^{2^{j-1}} - (1 - 0.5\lambda)^{2^j} \quad (6)$$

These polynomial *diffusion filters* are illustrated in Fig. 4 at different scales. They admit efficient implementation through repeated Laplacian multiplications, eliminating the need for explicit eigendecomposition. Given the input node features \mathbf{X} , the model aggregates the responses of all filters from both graphs as $\mathbf{H} = [\mathbf{H}_L \mid \mathbf{H}_H] \in \mathbb{R}^{N \times 2(J-1)F}$, where

$$\begin{aligned} \mathbf{H}_L &= [h_L^{(2)}(\mathbf{L}_{Ho})\mathbf{X} \mid h_L^{(3)}(\mathbf{L}_{Ho})\mathbf{X} \mid \cdots \mid h_L^{(J)}(\mathbf{L}_{Ho})\mathbf{X}], \\ \mathbf{H}_H &= [h_H^{(2)}(\mathbf{L}_{Ht})\mathbf{X} \mid h_H^{(3)}(\mathbf{L}_{Ht})\mathbf{X} \mid \cdots \mid h_H^{(J)}(\mathbf{L}_{Ht})\mathbf{X}]. \end{aligned}$$

Finally, a linear layer followed by a softmax function produces node-level class probabilities: $\hat{\mathbf{Y}} = \text{softmax}(\mathbf{H}\mathbf{W}_\phi)$, for $\mathbf{W}_\phi \in \mathbb{R}^{2(j-1)F \times C}$. Here, ϕ denotes the learnable weights of the linear layer in the classifier f_ϕ .

Structural loss. To explicitly enforce homophily and heterophily in the inferred graph, we define two *label-based structural losses*. The *homophilic structural loss* ℓ_{Ho} penalizes edges connecting dissimilar

labels, while the *heterophilic structural loss* ℓ_{Ht} penalizes edges connecting similar labels. Formally, they are defined as:

$$\ell_{\text{Ho}}(\mathcal{G}_{\text{Ho}}, \mathbf{Y}) = \sum_{(i,j) \in \mathcal{E}_f} w_{ij}^{(1)} (1 - \text{cosine}(\mathbf{y}_i, \mathbf{y}_j)), \quad (7)$$

$$\ell_{\text{Ht}}(\mathcal{G}_{\text{Ht}}, \mathbf{Y}) = \sum_{(i,j) \in \mathcal{E}_f} w_{ij}^{(2)} \text{cosine}(\mathbf{y}_i, \mathbf{y}_j), \quad (8)$$

where $\text{cosine}(\mathbf{y}_i, \mathbf{y}_j) = \mathbf{y}_i^\top \mathbf{y}_j / (\|\mathbf{y}_i\|_2 \|\mathbf{y}_j\|_2)$ is large for similar label vectors and $1 - \text{cosine}(\mathbf{y}_i, \mathbf{y}_j)$ is large for dissimilar ones. This ensures that a larger weight penalty is applied to the edge weights of nodes with the same class in the heterophilic mask, and a larger penalty is applied to the edge weights of nodes with different classes in the homophilic mask, thereby enforcing the heterophilic and homophilic structures for \mathcal{G}_{Ht} and \mathcal{G}_{Ho} , respectively. Recall the masking function introduced in (4), which outputs the weight $w_{ij}^{(k)} = \sigma(\Phi_{\theta_k}^\top(\mathbf{x}_i)\Phi_{\theta_k}(\mathbf{x}_j))$ for edge $(i, j) \in \mathcal{E}_f$. Since true labels are only available for a subset of nodes in the training set, we use the predicted softmax probabilities $\hat{\mathbf{Y}}$ in the structural loss. As we argue next in Proposition 1, the deviation between true and predicted label similarities is bounded by the prediction error, so the structural loss approaches the ideal label-based objective as predictions improve. Thus, the total loss ℓ in (3) is approximated as

$$\ell = \ell_{\text{CE}}(\hat{\mathbf{Y}}_{\text{train}}, \mathbf{Y}_{\text{train}}) + \ell_{\text{SL}}(\mathcal{G}_{\text{Ho}}, \mathcal{G}_{\text{Ht}}, \hat{\mathbf{Y}}),$$

where ℓ_{SL} follows the formulation in (3) with \mathbf{Y} replaced by $\hat{\mathbf{Y}}$. ℓ is minimized with respect to the learnable parameters θ_1 and θ_2 of the two masks, as well as the weights of the linear layer ϕ .

IV. THEORETICAL AND COMPLEXITY ANALYSES

Here, we derive stability guarantees to justify replacing labels with predictions and show that the filter bank output exhibits robustness to graph perturbations. Moreover, we discuss the complexity of FgGSL.

Structural loss stability. The structural loss ℓ_{SL} is computed using the predicted labels $\hat{\mathbf{Y}}$ in lieu of the ground-truth labels \mathbf{Y} . Using Proposition 1 below, we justify that this surrogate introduces only a controlled error, bounded by the discrepancy between \mathbf{Y} and $\hat{\mathbf{Y}}$.

Proposition 1. *Let C be the number of classes. Let $\hat{\mathbf{y}}_i$ and $\hat{\mathbf{y}}_j$ be the predicted class probabilities of the node vectors, and suppose $\|\mathbf{y}_i - \hat{\mathbf{y}}_i\|_2 \leq \epsilon_i$. Then*

$$|\text{cosine}(\mathbf{y}_i, \mathbf{y}_j) - \text{cosine}(\hat{\mathbf{y}}_i, \hat{\mathbf{y}}_j)| \leq 2\sqrt{C}(\epsilon_i + \epsilon_j).$$

The proof follows by writing the difference of the cosine similarities as an inner product of the unit vectors and then applying the Cauchy–Schwarz inequality.

Thus, as the classifier accuracy improves, the adopted structural loss approximation comes increasingly closer to the true label-based structural loss.

Filter bank stability under graph perturbations. Let \mathbf{L}^* denote the optimal graph Laplacian, and let $\hat{\mathbf{L}}$ be the recovered (possibly suboptimal) Laplacian used in the FgGSL filter bank. Following [10], we assume that the *operator distance modulo permutation* $\|\cdot\|_{\mathcal{P}}$ satisfies

$$\|\mathbf{L}^* - \hat{\mathbf{L}}\|_{\mathcal{P}} = \|\mathbf{E}\| \leq \epsilon,$$

where \mathbf{E} is the absolute perturbation modulo permutation. For a Lipschitz graph filter $h_j(\cdot)$ with Lipschitz constant K_h , [10, Theorem 1] establishes the following bound on the filter outputs:

$$\|h_j(\mathbf{L}^*) - h_j(\hat{\mathbf{L}})\|_{\mathcal{P}} \leq 2K_h(1 + \delta\sqrt{N})\epsilon + \mathcal{O}(\epsilon^2),$$

where $\delta := (\|\mathbf{U} - \mathbf{V}\|_2 + 1)^2 - 1$ quantifies the eigenvector misalignment between the Laplacian $\mathbf{L}^* = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^\top$ and the error matrix $\mathbf{E} = \mathbf{V}\mathbf{\Lambda}_E\mathbf{V}^\top$.

Applying this result to our multiscale filter bank in (5) yields

$$\|h_j(\mathbf{L}^*) - h_j(\hat{\mathbf{L}})\|_{\mathcal{P}} \leq 2^{(j-1)}(1 + \delta\sqrt{N})\epsilon + \mathcal{O}(\epsilon^2).$$

This indicates that the filter banks remain stable under small structural perturbations of the graph, and a near-optimal graph structure will have a minimal effect on the node embeddings generated by FgGSL.

Computational complexity. FgGSL consists of two major components: (i) mask construction; and (ii) graph filtering. Given a graph with N nodes, the initial fully-connected candidate graph contains up to $\frac{N(N-1)}{2}$ edges. Computing the masking weights defined as inner products of transformed node features requires $\mathcal{O}(N^2)$ operations. After the masked graph is obtained, applying a single filter of the form $h_j(\mathbf{L})\mathbf{X}$ incurs a cost of $\mathcal{O}(N^2f)$, where f is the feature dimension. Constructing the set of 2-to- J filters, i.e., $[h_2(\mathbf{L}), h_3(\mathbf{L}), \dots, h_J(\mathbf{L})]$, requires $\mathcal{O}(JN^3)$ time due to repeated multiplication with the graph Laplacian. Overall, the dominant cost arises from building the multi-scale filter bank, yielding a total complexity of $\mathcal{O}(JN^3 + N^2f)$.

V. NUMERICAL EXPERIMENTS

In this section, we evaluate the effectiveness of the proposed FgGSL framework on node classification tasks over 6 heterophilic graph datasets. We compare our approach against a comprehensive set of baselines: Geom-GCN [2], a geometry-aware architecture that preserves structural and positional relations between nodes; GAT [11], an attention-based GNN that learns edge-wise importance weights; GraphSAGE [12], an inductive neighborhood-sampling model with feature aggregation; a graph-agnostic MLP; H2GCN [6], a heterophily-aware model that separates ego and neighbor information while leveraging higher-order neighborhoods; FAGCN [5], which incorporates high-frequency information in GCN layers; MixHop [4], which aggregates and mixes multi-hop neighborhood information in parallel; and SG-GCN [8], which generates multiple parallel graph structures based on different score functions and trains separate GNNs to produce the final node embeddings. The performance metrics are reported as mean and standard deviation of node classification accuracy over ten standard train–val–test splits.

We further conduct an ablation study to assess the contribution of each module in our framework. Three variants are considered: FgGSL (NM), FgGSL (FBL), and FgGSL (FBH). The FgGSL (NM) variant removes the masking mechanism and applies the two parallel

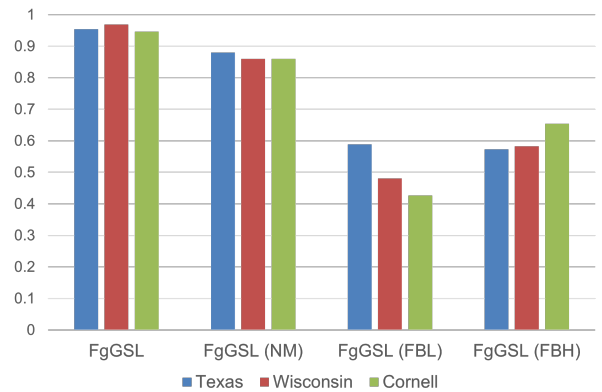


Fig. 5: Mean accuracy of different variants of the model.

TABLE I: Performance Comparison on Heterophilic Graph Datasets.

Model	Texas (0.88)	Wisconsin (0.79)	Cornell (0.87)	Squirrel (0.77)	Actor (0.78)	Chameleon (0.76)
FgGSL	0.94 ± 0.08	0.96 ± 0.05	0.94 ± 0.08	0.58 ± 0.09	0.41 ± 0.02	0.79 ± 0.09
SAGE	0.74 ± 0.08	0.74 ± 0.08	0.69 ± 0.05	0.37 ± 0.02	0.34 ± 0.01	0.50 ± 0.01
GAT	0.52 ± 0.06	0.49 ± 0.04	0.61 ± 0.05	0.40 ± 0.01	0.27 ± 0.01	0.60 ± 0.02
MLP	0.79 ± 0.04	0.85 ± 0.03	0.75 ± 0.02	0.35 ± 0.02	0.35 ± 0.01	0.50 ± 0.02
H2GCN	0.80 ± 0.05	0.84 ± 0.05	0.70 ± 0.05	0.59 ± 0.01	0.35 ± 0.01	0.69 ± 0.01
Geom-GNN	0.78 ± 0.07	0.80 ± 0.06	0.61 ± 0.08	0.56 ± 0.02	0.35 ± 0.01	0.65 ± 0.02
MixHop	0.81 ± 0.09	0.83 ± 0.08	0.78 ± 0.09	0.35 ± 0.03	0.34 ± 0.01	0.53 ± 0.02
SG-GCN	0.83 ± 0.01	0.83 ± 0.01	0.72 ± 0.01	0.60 ± 0.02	0.36 ± 0.01	0.67 ± 0.03
FAGCN	0.83 ± 0.01	0.82 ± 0.01	0.71 ± 0.01	0.31 ± 0.02	0.35 ± 0.01	0.46 ± 0.03




Fig. 6: Distribution of cosine similarities of node embeddings for all node pairs in the Texas (left) and Cornell (right) heterophilic datasets. Notice how node pairs with the same labels have higher cosine similarity compared to those with different labels, demonstrating the effectiveness of the embeddings. This is to be compared with the overlapping distributions in Fig. 1.

filter banks directly on the given graph. The FgGSL (FBL) variant retains the masking function but employs only the low-pass filter bank to process the nodal features. Conversely, FgGSL (FBH) uses masking followed by the high-pass filter bank only. This way, we isolate the masking mechanism and spectral filtering components.

Datasets. We conduct experiments on six widely used heterophilic benchmark datasets [2]: Texas (heterophily ratio: 0.88), Wisconsin (0.79), Cornell (0.87), Squirrel (0.77), Chameleon (0.76), and Actor (0.78). The first three datasets originate from the WebKB domain and comprise webpage graphs with strong label dissimilarity. In contrast, Squirrel and Chameleon represent large Wikipedia page networks characterized by high feature dimensionality and noisy inter-class edges. The Actor dataset contains co-occurrence relations between actors in film contexts, exhibiting pronounced heterophily.

Results. Across all six heterophilic graph datasets, FgGSL consistently achieves top or competitive accuracy. As shown in Table I, FgGSL outperforms or matches the performance of existing baselines on every dataset. An ablation study, shown in Fig. 5, demonstrates that the full model performs best across multiple datasets. The accuracy drop in FgGSL (NM) highlights the importance of feature-guided masking, while the weaker performance of the single-filter variants underscores the benefit of combining both filter banks.

To further evaluate the quality of the learned node embeddings, we analyze the cosine similarity distributions of intra-class and inter-class node pairs. Unlike raw features (cf. Fig. 1), the embeddings from FgGSL exhibit a clear separation between the two types of node pairs, as illustrated in Fig. 6. This separation demonstrates the model’s effectiveness in capturing heterophilic structures.

VI. CONCLUSIONS

In this work, we introduced FgGSL, a frequency-guided framework for joint graph structure learning and node classification on heterophilic graphs. By learning feature-driven masks for homophilic and heterophilic graphs (thereby learning two graphs) and combining low-pass and high-pass spectral filters, FgGSL recovers both

homophilic and heterophilic relations that are informative for the downstream task. Our label-based structural loss enables supervised, task-driven graph inference and is supported by stability guarantees for both the loss and the filter bank. Empirical evaluations across six heterophilic benchmarks show that FgGSL consistently outperforms existing GNNs and rewiring approaches, while maintaining an interpretable and theoretically grounded design.

REFERENCES

- [1] W. Hamilton, *Graph Representation Learning*, ser. Synthesis Lectures on Artificial Intelligence and Machine Learning. Springer Nature Switzerland, 2022. [Online]. Available: <https://books.google.co.in/books?id=6YRyEAAAQBAJ>
- [2] H. Pei, B. Wei, K. C.-C. Chang, Y. Lei, and B. Yang, “Geom-GCN: Geometric Graph Convolutional Networks,” Feb. 2020, arXiv:2002.05287 [cs]. [Online]. Available: <http://arxiv.org/abs/2002.05287>
- [3] S. Luan, C. Hua, Q. Lu, L. Ma, L. Wu, X. Wang, M. Xu, X.-W. Chang, D. Precup, R. Ying, S. Z. Li, J. Tang, G. Wolf, and S. Jegelka, “The Heterophilic Graph Learning Handbook: Benchmarks, Models, Theoretical Analysis, Applications and Challenges,” Jul. 2024, arXiv:2407.09618 [cs]. [Online]. Available: <http://arxiv.org/abs/2407.09618>
- [4] S. Abu-El-Hajja, B. Perozzi, A. Kapoor, N. Alipourfard, K. Lerman, H. Harutyunyan, G. V. Steeg, and A. Galstyan, “MixHop: Higher-Order Graph Convolutional Architectures via Sparsified Neighborhood Mixing,” in *Proceedings of the 36th International Conference on Machine Learning*. PMLR, May 2019, pp. 21–29, ISSN: 2640-3498. [Online]. Available: <https://proceedings.mlr.press/v97/abu-el-hajja19a.html>
- [5] D. Bo, X. Wang, C. Shi, and H. Shen, “Beyond Low-frequency Information in Graph Convolutional Networks,” Jan. 2021, arXiv:2101.00797 [cs]. [Online]. Available: <http://arxiv.org/abs/2101.00797>
- [6] J. Zhu, Y. Yan, L. Zhao, M. Heimann, L. Akoglu, and D. Koutra, “Beyond Homophily in Graph Neural Networks: Current Limitations and Effective Designs,” Oct. 2020, arXiv:2006.11468 [cs]. [Online]. Available: <http://arxiv.org/abs/2006.11468>
- [7] W. Bi, L. Du, Q. Fu, Y. Wang, S. Han, and D. Zhang, “Make Heterophily Graphs Better Fit GNN: A Graph Rewiring Approach,” Sep. 2022, arXiv:2209.08264 [cs]. [Online]. Available: <http://arxiv.org/abs/2209.08264>
- [8] V. M. Tenorio, M. Navarro, S. Rey, S. Segarra, and A. G. Marques, “Adapting to Heterophilic Graph Data with Structure-Guided Neighbor Discovery,” Jun. 2025, arXiv:2506.08871 [cs] version: 1. [Online]. Available: <http://arxiv.org/abs/2506.08871>
- [9] X. Dong, D. Thanou, M. Rabbat, and P. Frossard, “Learning Graphs From Data: A Signal Representation Perspective,” *IEEE Signal Processing Magazine*, vol. 36, no. 3, pp. 44–63, May 2019. [Online]. Available: <https://ieeexplore.ieee.org/document/8700665>
- [10] F. Gama, J. Bruna, and A. Ribeiro, “Stability Properties of Graph Neural Networks,” *IEEE Transactions on Signal Processing*, vol. 68, pp. 5680–5695, 2020. [Online]. Available: <https://ieeexplore.ieee.org/document/9206091>
- [11] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio, “Graph Attention Networks,” Feb. 2018, arXiv:1710.10903 [stat]. [Online]. Available: <http://arxiv.org/abs/1710.10903>
- [12] W. L. Hamilton, R. Ying, and J. Leskovec, “Inductive Representation Learning on Large Graphs,” Sep. 2018, arXiv:1706.02216 [cs]. [Online]. Available: <http://arxiv.org/abs/1706.02216>