



Characterizing the Influence of Topology on Graph Learning Tasks

Kailong Wu¹, Yule Xie¹, Jiaxin Ding¹✉, Yuxiang Ren², Luoyi Fu¹, Xinbing Wang¹, and Chenghu Zhou¹

¹ Shanghai Jiao Tong University Shenzhen, China
{1473686097, xyl-alter, jiaxinding, yiluofu, xwang8}@sjtu.edu.cn,
zhouchsjtu@gmail.com

² 2012 Laboratories, Huawei Technologies Shenzhen, China
renyuxiang1@huawei.com

Abstract. Graph neural networks (GNN) have achieved remarkable success in a wide range of tasks by encoding features combined with topology to create effective representations. However, the fundamental problem of understanding and analyzing how graph topology influences the performance of learning models on downstream tasks has not yet been well understood. In this paper, we propose a metric, TopoInf, which characterizes the influence of graph topology by measuring the level of compatibility between the topological information of graph data and downstream task objectives. We provide analysis based on the decoupled GNNs on the contextual stochastic block model to demonstrate the effectiveness of the metric. Through extensive experiments, we demonstrate that TopoInf is an effective metric for measuring topological influence on corresponding tasks and can be further leveraged to enhance graph learning.

Keywords: GNNs, graph topology, graph rewiring

1 Introduction

Graph neural networks (GNNs) have emerged as state-of-the-art models to learn graph representation by the message-passing mechanism over graph topology for downstream tasks, such as node classification, link prediction, etc. [12, 3, 18].

Despite their success, GNNs are vulnerable to the compatibility between graph topology and graph tasks [14, 9]. In essence, the effectiveness of the message-passing mechanism and, by extension, the performance of GNNs depends on the assumption that the way information is propagated and aggregated in the graph neural networks should be conducive to the downstream tasks. First, if the underlying motivation of graph topology, why two nodes get connected, has no relation to the downstream task, such task-irrelevant topology can hurt the information aggregation per se. For example, in a random graph that lacks meaningful clusters, graph learning becomes impossible in community detection tasks. Furthermore, GNNs may even be outperformed by multilayer perceptrons

(MLP) on graphs exhibiting heterophily [31, 32]. Second, even if the motivation for connections is compatible with tasks, graphs in real-world applications are usually inherent with “noisy” edges or incompleteness, due to error-prone data collection. This topological noise can also affect the effectiveness of message passing. Third, the vulnerability of GNNs to compatibility varies across models due to differences in their message-passing mechanisms. For example, GCN [12] may not be as effective on heterophilic graphs. However, H2GCN [32] is able to mitigate the heterophilic problem, with specially designed message passing. Furthermore, even the same model with different hyperparameters of network layers can also exhibit different vulnerabilities, since the reception fields of the neighborhoods are different. All this, there is a natural question to ask: *What is the metric for the compatibility between graph topology and graph learning tasks?*

This question, which is essential for choosing a graph learning model to deploy over a graph for downstream tasks, has not yet been systematically characterized. Existing work focuses on measuring homophily [17, 32, 24] or edge signal-to-noise ratio [9], based on the discrepancy of nodes and their neighbors in features or labels without considering graph learning models. Therefore, the results of these metrics can be an indicator of the performance of GCN-like models but cannot be generalized to more recently developed GNN models and guide topology modification to improve the performance of these models. From the above analysis, we would like to investigate the compatibility between topology and tasks associated with graph learning models globally, and locally, the influence from each edge, either positive or negative, can be used to adjust the topology to increase or decrease the performance. Answering this question can help us better understand graph learning, increase the interpretability of learning models, and shed light on improving model performance.

We first propose a metric for model-dependent compatibility between topology and graph tasks, measured by the difference between the “ideal” results, labels, and the result of the models performing on the “ideal” features. Thereafter, we introduce a metric TopoInf, to locally characterize the influence of each edge on the overall compatibility, by evaluating the change of such compatibility after topology perturbation. We provide motivating analysis based on the decoupled GNNs [13, 8] on cSBMs [7] to validate the effectiveness of the metric and demonstrate that TopoInf is an effective metric through extensive experiments.

Our main contributions can be summarized as follows:

- To the best of our knowledge, we are the first to measure compatibility between graph topology and learning tasks associated with graph models.
- We propose a new metric, TopoInf, to measure the influence of edges on the performance of GNN models in node classification tasks, and conduct extensive experiments to validate the effectiveness.
- The proposed TopoInf can be leveraged to improve performance by modifying the topology on different GNN models, and such a scheme can be applied further to different scenarios.

2 Preliminaries

Notations. We denote an undirected graph as $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{A})$, where $\mathcal{V}$ is the node set with cardinality $|\mathcal{V}| = n$, $\mathcal{E}$ is the edge set without self-loop, and $\mathbf{A} \in \mathbb{R}^{n \times n}$ is the symmetric adjacency matrix with $\mathbf{A}_{i,j} = 1$ when $e_{ij} \in \mathcal{E}$ otherwise $\mathbf{A}_{i,j} = 0$. $\mathbf{D}$ denotes the diagonal degree matrix of $\mathcal{G}$ where $\mathbf{D}_{i,i} = d_i$, the degree of node v_i . We use $\tilde{\mathbf{A}} = \mathbf{A} + \mathbf{I}$ to represent the adjacency matrix with self-loops and $\tilde{\mathbf{D}} = \mathbf{D} + \mathbf{I}$. The symmetric normalized adjacency matrix is $\hat{\mathbf{A}} = \mathbf{D}^{-1/2} \tilde{\mathbf{A}} \tilde{\mathbf{D}}^{-1/2}$. For node classification with c classes, $\mathbf{L} \in \mathbb{R}^{n \times c}$ denotes the label matrix, whose i^{th} row represents the one-hot encoding of node v_i , while $\mathbf{X} \in \mathbb{R}^{n \times d}$ is the feature matrix, whose i^{th} row $\mathbf{X}_{i,:}$ represents a d -dimensional feature vector of node v_i .

Graph Filters. Recent studies show that GNN models, such as ChebNet [6], APPNP [13] and GPRGNN [4], can be viewed as operations of polynomial graph spectral filters. Specifically, the effect of such GNN models on a graph signal $\mathbf{x} \in \mathbb{R}^{n \times 1}$, can be formulated as graph spectral filtering operation $\mathbf{f}(\mathbf{A})$ based on adjacency matrix $\mathbf{A}$, such that $\mathbf{f}(\mathbf{A})\mathbf{x} = \sum_{k=0}^K \gamma_k \hat{\mathbf{A}}^k \mathbf{x}$, where $\hat{\mathbf{A}}$ is the normalized adjacency matrix, K is the order of the graph filter and γ_k 's are the weights, which can be assigned [26, 13] or learned [4, 11]. The graph filter can be extracted as the effect of topology associated with GNN models.

Decoupled GNN. Recent works [13, 29, 8] show that neighborhood aggregation and feature transformation can be decoupled and formulate the graph learning tasks as

$$\hat{\mathbf{L}} = \text{softmax}(\mathbf{f}(\mathbf{A})\mathbf{g}_\theta(\mathbf{X})) \quad (1)$$

where the prediction $\hat{\mathbf{L}}$ can be obtained by operating on graph feature matrix $\mathbf{X}$, through the feature transformation function $\mathbf{g}_\theta(\cdot)$ with learnable parameters, thereafter applying the graph filter $\mathbf{f}(\mathbf{A})$ acting as the neighborhood aggregation, and finally passing through a softmax layer.

3 Methodology

In this section, we present our methodology to study the influence of topology on graph learning tasks. We evaluate the global compatibility between graph topology and tasks, and propose our metric, TopoInf.

3.1 Matching graph topology and tasks

We first relate the graph topology to the graph tasks. Thanks to the graph filters in the preliminaries, the influence of the graph topology on a GNN model can be simplified and approximated as a graph filter $\mathbf{f}(\mathbf{A})$. $\mathbf{f}(\mathbf{A})$ can be viewed as a graph filter function $\mathbf{f}(\cdot) : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^{n \times n}$ applied to the adjacency matrix $\mathbf{A}$ of the observed graph. With the decoupling analysis of GNN, the prediction is obtained by $\mathbf{f}(\mathbf{A})$ working on prediction results $\mathbf{g}_\theta(\mathbf{X})$ based on feature matrix $\mathbf{X}$ through graph learning models [5, 29, 8], where $\mathbf{g}_\theta(\mathbf{X})$ on $\mathbf{X}$ can be viewed

as extracting task-related information in $\mathbf{X}$ and ideally, one of the best predictions on $\mathbf{X}$ could be the label matrix $\mathbf{L}$. Therefore, we replace the functions with $\mathbf{f}(\mathbf{A})\mathbf{g}_\theta(\mathbf{X})$ in Equation 1 to be $\mathbf{f}(\mathbf{A})\mathbf{L}$, which simplifies the analysis from features, makes us focus on the labels and graph tasks, and provides an “upper bound” of the predicted results in the ideal situation. In the next section, we will provide a motivation analysis for this setting.

Further, instead of applying the softmax function, we perform row-normalization [9], which normalizes the summation of entries on the same row to be one, on $\mathbf{f}(\mathbf{A})\mathbf{L}$ as the prediction as used in previous works [16, 9]. The row-normalized matrix $\mathbf{f}(\mathbf{A})\mathbf{L}$ provides a summary of the edge statistics at the individual node level, as each row of $\mathbf{f}(\mathbf{A})\mathbf{L}$ represents the label neighborhood distribution of the node, which has direct connections to GNN learning [15]. As $\text{RowNorm}(\mathbf{f}(\mathbf{A})\mathbf{L}) = \text{RowNorm}(\mathbf{f}(\mathbf{A}))\mathbf{L}$, we perform row-normalization on graph filter $\mathbf{f}(\mathbf{A})$, and $\mathbf{f}(\mathbf{A})$ is row-normalized in our subsequent analysis.

Therefore, to quantify the compatibility between graph topology and graph tasks, we propose measuring the similarity between the ideal prediction $\mathbf{L}$ and the graph filter that performs the ideal feature $\mathbf{f}(\mathbf{A})\mathbf{L}$. The more similar $\mathbf{L}$ and $\mathbf{f}(\mathbf{A})\mathbf{L}$ is ($\mathbf{f}(\mathbf{A})\mathbf{L} \sim \mathbf{L}$ for short), the better the graph topology matches the graph tasks, which we provide a motivating example with mild assumptions to validate in the next section. This can be understood through various scenarios:

- In graph filtering, the fact that $\mathbf{f}(\mathbf{A})\mathbf{L} \sim \mathbf{L}$ indicates that signals in $\mathbf{L}$, whether low or high frequency, are well preserved after being filtered by $\mathbf{f}(\mathbf{A})$. This preservation of information suggests a good match between task labels and the graph topology.
- Considering the concept of the homophily/heterophily, when $\mathbf{f}(\mathbf{A})\mathbf{L} \sim \mathbf{L}$ means linked nodes described by $\mathbf{f}(\mathbf{A})$ likely belong to the same class or have similar characteristics, which is aligned with the homophily principle.
- In terms of Dirichlet energy, if $\mathbf{f}(\mathbf{A})\mathbf{L} \sim \mathbf{L}$ means low Dirichlet energy of $\mathbf{L}$ w.r.t. the (augmented) normalized Laplacian derived from $\mathbf{f}(\mathbf{A})$. This indicates the smoothness of $\mathbf{L}$ on $\mathbf{f}(\mathbf{A})$, suggesting a good match between the graph topology and tasks.

All above, we define $\mathcal{I}(\mathbf{A})$ to quantify the degree of compatibility between graph topology and graph tasks,

$$\mathcal{I}(\mathbf{A}) = \mathcal{S}(\mathbf{L} \parallel \mathbf{f}(\mathbf{A})\mathbf{L}),$$

where $\mathcal{S}(\cdot)$ is a similarity measurement, which can be chosen according to the requirements of the problem, for example, similarity induced by the inner product or Euclidean distance.

To find a graph topology that matches the graph tasks well, our goal is to maximize $\mathcal{I}(\mathbf{A})$ over $\mathbf{A}$. However, $\mathcal{I}(\mathbf{A})$ would be trivially maximized with $\mathbf{f}(\mathbf{A}) = \mathbf{I}$, where the graph topology is completely removed. To address this issue, we introduce a regularization function $\mathcal{R}(\cdot)$, which penalizes $\mathcal{I}(\mathbf{A})$ if $\mathbf{f}(\mathbf{A})$ is close to the identity matrix, to minimize the modification of $\mathbf{f}(\mathbf{A})$. We also

demonstrate such regularization from the perspective of graph denoising in the next section. Therefore, the optimization problem is formulated as

$$\max_{\mathbf{A}} \mathcal{C}(\mathbf{A}) = \mathcal{I}(\mathbf{A}) - \lambda \mathcal{R}(\mathbf{A}) \quad (2)$$

where λ is a hyperparameter used to balance the trade-off between $\mathcal{I}(\mathbf{A})$ and $\mathcal{R}(\mathbf{A})$. However, there are $2^{n \times n}$ possible states of $\mathbf{A}$, making the optimization of $\mathcal{C}(\mathbf{A})$ computationally intractable due to its nondifferentiability and combinatorial nature. Moreover, the optimization does not consider the original topology. Therefore, instead of optimizing the problem globally, a more practical and effective approach is to start with the originally observed graph and locally optimize it, as we will present in the next subsection.

Now we zoom in the general compatibility to the node level: for every node v_i , we use the node level similarity function $\mathcal{S}_{\mathbf{v}}(\cdot)$ to measure the similarity of its one-hot label vector $\mathbf{L}_{i,:}$ and its normalized soft label vector $\bar{\mathbf{L}}_{i,:}$ filtered by $\mathbf{f}(\mathbf{A})$, and node level regularization function $\mathcal{R}_{\mathbf{v}}(\cdot)$ to measure the regularization of its connectivity to other nodes, specifically, the i^{th} row of $\mathbf{f}(\mathbf{A})$. Let $\mathcal{C}_{\mathbf{A}}(v_i)$ denote the compatibility between the topology of node v_i and graph tasks. The overall compatibility metric becomes the following

$$\mathcal{C}(\mathbf{A}) = \sum_{v_i \in \mathcal{V}_t} \mathcal{C}_{\mathbf{A}}(v_i) = \sum_{v_i \in \mathcal{V}_t} \mathcal{I}_{\mathbf{A}}(v_i) - \lambda \mathcal{R}_{\mathbf{A}}(v_i), \quad (3)$$

where $\mathcal{I}_{\mathbf{A}}(v_i) = \mathcal{S}_{\mathbf{v}}(\mathbf{L}_{i,:} \| \bar{\mathbf{L}}_{i,:})$ and $\mathcal{V}_t \subseteq \mathcal{V}$ is the node set. Notice that $\mathcal{V}_t$ can be not only the whole node set but also be chosen flexibly as the node set whose compatibility we care about most in correspondence to the circumstances. For example, in adversarial attack tasks, $\mathcal{V}_t$ can be a collection of targeted nodes to be attacked; in graph node classification tasks, $\mathcal{V}_t$ can be the nodes to be predicted. In this work, we choose the inner product as the similarity metric, where $\mathcal{I}_{\mathbf{A}}(v_i)$ can be evaluated as $\mathcal{I}_{\mathbf{A}}(v_i) = \bar{\mathbf{L}}_{i,c_i}$, and c_i is the true label of v_i , and the reciprocal of degree as the regularizer, where $\mathcal{R}_{\mathbf{A}}(v_i)$ can be evaluated as $\mathcal{R}_{\mathbf{A}}(v_i) = 1/d_i$, and d_i is the degree of node v_i .

3.2 TopoInf: measuring the influence of graph topology on tasks

From the above analysis, we can evaluate how well the overall graph topology and learning tasks are in accordance with $\mathcal{C}(\mathbf{A})$. We are interested in optimizing $\mathcal{C}(\mathbf{A})$, which can improve the compatibility between graph topology and tasks, and, by extension, improve the performance of the learning model. We are also interested in characterizing the influence of modifying part of topology on the degree of task compatibility, which can provide more interpretability of graph learning and meanwhile guide the topology modification for better task performance. To maximize $\mathcal{C}(\mathbf{A})$, we can take the “derivative” of $\mathcal{C}(\mathbf{A})$ by obtaining the change of $\mathcal{C}(\mathbf{A})$ with local topology perturbation. Here we focus on the topology of edges and the same analysis can be simply extended to nodes and substructures. The influence of edge e_{ij} can be measured by the difference between $\mathcal{C}(\mathbf{A})$

on original (normalized) adjacency matrix $\mathbf{A}$ and $\mathcal{C}(\mathbf{A}')$ on the modified (normalized) adjacency matrix $\mathbf{A}'$ obtained by removing edge e_{ij} . Therefore, we define such topology influence as TopoInf of edge e_{ij} , denoted as $\nabla\mathcal{C}_{\mathbf{A}}(e_{ij})$, is given by

$$\nabla\mathcal{C}_{\mathbf{A}}(e_{ij}) = \mathcal{C}(\mathbf{A}') - \mathcal{C}(\mathbf{A}) = \sum_{v_k \in \mathcal{V}} \mathcal{C}_{\mathbf{A}'}(v_k) - \mathcal{C}_{\mathbf{A}}(v_k). \quad (4)$$

The sign of **TopoInf** reflects the positive or negative influence of removing the edge on the matching of topology and tasks, which also indicates that removing the edge can increase or decrease the model performance. Remark that edges with positive TopoInf correspond to edges with a “negative” effect on the model performance, such that removing those can bring a positive influence, and vice versa. The absolute value of **TopoInf** measures the magnitude of influence: a higher absolute value means a higher influence.

It should be noted that to compute **TopoInf** of an edge, we do not need to recompute $\mathbf{f}(\mathbf{A}')\mathbf{L}$, which is computationally expensive. The difference between $\mathcal{I}(\mathbf{A}')$ and $\mathcal{I}(\mathbf{A})$ after removing e_{ij} is limited within K hop neighborhood of v_i and v_j considering a K -order filter $\mathbf{f}(\mathbf{A})$. Therefore, we only need to compute the difference of node influence on the neighborhood affected by the edge removal in Equation 3.

Remark that we do not assume that we have all the true labels. Using all labels is for the convenience of analysis to demonstrate the effectiveness of the metric. In practical graph tasks where not all labels are available, we can use pseudo labels obtained by MLP or GNN models as replacements for true labels and compute estimated TopoInf based on pseudo labels. This will introduce errors, but can still be effective which we will show in our experiment.

Moreover, due to the presence of nonlinearity in GNN models, the extraction of $\mathbf{f}(\mathbf{A})$ can be challenging for graph models, and an approximation of $\mathbf{f}(\mathbf{A})$ is required. We assume that $\mathbf{f}(\mathbf{A})$ can be obtained or approximated for the graph models, and a better approximation can improve the accuracy of the influence metric. Here, we present our approach to obtaining, approximating and learning $\mathbf{f}(\mathbf{A})$ for several representative GNN models.

- **Obtained graph filters.** For SGC [26], S2GC [30], APPNP [13], and other decoupled GNNs, $\mathbf{f}(\mathbf{A})$ can be directly obtained by Equation 1.
- **Approximated graph filters.** For GCN [12] and other GCN-like models, the presence of nonlinearities introduces additional complexities to $\mathbf{f}(\mathbf{A})$, as it becomes dependent not only on the observed graph topology but also on the learnable parameters and the activation function. In this study, we adopt an approximate approach to estimate $\mathbf{f}(\mathbf{A})$ by removing the nonlinear activation to simplify the analysis, as works in [26, 15]. For GCNII [3] and other GNNs with residual connections between layers, we can employ a similar approximate approach as in GCN to estimate $\mathbf{f}(\mathbf{A})$ by removing the nonlinear normalization. This approximation without nonlinearity, can still capture the effects of graph learning models performing on the topology, as we will show in our experiments.

- **Learned graph filters.** For GPRGNN [4], ChebNet [6], and other GNNs driven by learning filters, the filter weights are learned based on backpropagation, where we are not able to obtain the graph filter $\mathbf{f}(\mathbf{A})$ before training. In this study, we adopt two approximate approaches to deal with this problem. One approach is to train the filter-based GNN model, obtain the trained filter weights, and then apply the learned graph filter as $\mathbf{f}(\mathbf{A})$. Another approach is to use fixed filter weights or predefined filters as the graph filter $\mathbf{f}(\mathbf{A})$, such as $\hat{\mathbf{A}}^K$. This approach ignores the effect of the learned filter weights in the GNN model but can still preserve the order information and other prior knowledge from the GNN model. Moreover, for GAT [25] and other attention-based GNNs, unlike GNNs driven by learning filters, the learned filters are implicitly determined by the weights of the neural network and the input features, and the learned filters may change due to the randomness during initialization and training. In this case, we apply only the second approach with a predefined fixed filter as an approximation.

Table 1: **Approximation approach of $\mathbf{f}(\mathbf{A})$ for representative GNN models.** Here we take K-layer GNN models as examples. α is predefined hyper-parameter and γ is learnable parameter. $\mathbf{H}^{(0)} = \mathbf{g}_\theta(\mathbf{X})$.

Model	Neighbor Aggregation	Filter	Type
SGC	$\hat{\mathbf{L}} = \text{softmax}(\hat{\mathbf{A}}^K \mathbf{X} \mathbf{W})$	$\mathbf{f}(\mathbf{A}) = \hat{\mathbf{A}}^K$	obtained
S2GC	$\hat{\mathbf{L}} = \text{softmax}\left(\frac{1}{K} \sum_{k=1}^K \left((1-\alpha)\hat{\mathbf{A}}^k + \alpha I\right) \mathbf{X} \mathbf{W}\right)$	$\mathbf{f}(\mathbf{A}) = \frac{1}{K} \sum_{k=1}^K (1-\alpha)\hat{\mathbf{A}}^k + \alpha I$	obtained
APPNP	$\mathbf{H}^{(\ell+1)} = (1-\alpha)\hat{\mathbf{A}}\mathbf{H}^{(\ell)} + \alpha\mathbf{H}^{(0)}$	$\mathbf{f}(\mathbf{A}) = (1-\alpha)\hat{\mathbf{A}}^K + \sum_{k=0}^{K-1} \alpha(1-\alpha)^k \hat{\mathbf{A}}^k$	obtained
GCN	$\mathbf{H}^{(\ell+1)} = \sigma\left(\hat{\mathbf{A}}\mathbf{H}^{(\ell)}\mathbf{W}^{(\ell)}\right)$	$\mathbf{f}(\mathbf{A}) = \hat{\mathbf{A}}^K$	approximated
GCNII	$\mathbf{H}^{(\ell+1)} = \sigma\left(\left((1-\alpha)\hat{\mathbf{A}}\mathbf{H}^{(\ell)} + \alpha\mathbf{H}^{(0)}\right)\mathbf{W}^{(\ell)}\right)$	$\mathbf{f}(\mathbf{A}) = (1-\alpha)\hat{\mathbf{A}}^K + \sum_{k=0}^{K-1} \alpha(1-\alpha)^k \hat{\mathbf{A}}^k$	approximated
GPRGNN	$\hat{\mathbf{L}} = \text{softmax}\left(\sum_{k=0}^K \gamma_k \hat{\mathbf{A}}^k \mathbf{H}^{(0)}\right)$	$\mathbf{f}(\mathbf{A}) = \sum_{k=0}^K \gamma_k \hat{\mathbf{A}}^k$	learned

4 Motivation

In this section, we provide a theoretical analysis and motivating example to validate our methodology and increase its comprehensibility. Specifically, we provide alternative perspectives to explain the meaning and validate the necessity of $\mathcal{I}(\mathbf{A})$ and $\mathcal{R}(\mathbf{A})$ in $\mathcal{C}(\mathbf{A})$.

We use contextual stochastic block model (cSBM) [7] for analysis, which is a widely used model for complex graph structures that integrates contextual features and graph topology [1, 9, 7]. In cSBM, node features are generated by the spiked covariance model, which follows a Gaussian distribution, the mean of which depends on the community assignment. The underlying assumption is that the node features are “informative” and can be perceived as label embedding vectors with Gaussian noise. The graph topology is generated by SBMs, resulting in communities of nodes connected by certain edge densities.

Definition 1. (Contextual Stochastic Block Model, cSBM). $\mathcal{G}$ has n nodes belonging to c communities, with intra/inter-community edge probabilities of p and q . The feature matrix is $\mathbf{X} = \mathbf{F} + \mathbf{N}$, where $\mathbf{N}_{ij} \sim \mathcal{N}(0; \sigma^2)$ are i.i.d. Gaussian noise, $\mathbf{F}_{i,:} = \boldsymbol{\mu}_{c_i}^\top$ is the d dimensional feature vector for the center

of the community c_i , which is the community v_i belongs to. Therefore, $\mathbf{F} = \mathbf{L}\boldsymbol{\mu}$, where $\mathbf{L}$ is the matrix of one-hot labels that denotes the community to which each node belongs, and $\boldsymbol{\mu} = (\boldsymbol{\mu}_1, \boldsymbol{\mu}_2, \dots, \boldsymbol{\mu}_c)^\top \in \mathbb{R}^{c \times d}$.

4.1 Theoretical Analysis

Here we present a theoretical analysis to validate our TopoInf and enhance its comprehensibility. According to the definition of cSBMs [7], the feature matrix $\mathbf{F} = \mathbf{L}\boldsymbol{\mu}$ is directly related to the prediction $\mathbf{L}$ and contains sufficient information for the graph learning task, which is also discussed and empirically verified in [16]. In this part, we follow the setting in [16] as a motivating example, and treat $\mathbf{F}$ as the true signals and $\mathbf{f}(\mathbf{A})\mathbf{X}$ as the prediction, and evaluate the effect of the low-pass filter by measuring the difference between the filtered feature matrix $\mathbf{f}(\mathbf{A})\mathbf{X}$ and the true feature matrix $\mathbf{F}$. To analyze the difference, we focus on the low-pass filter, which is true for most GCN-like models [12, 26, 13, 16], use $\|\mathbf{f}(\mathbf{A})\mathbf{X} - \mathbf{F}\|$ as the prediction error introduced by the learning model and the topology, and consequently establish the following theorem.

Theorem 1. *For $0 < \delta < 1$, with probability at least $1 - \delta$, we have*

$$\|\mathbf{f}(\mathbf{A})\mathbf{X} - \mathbf{F}\| \leq c_1 \|\mathbf{f}(\mathbf{A})\mathbf{L} - \mathbf{L}\| + c_2 \|\mathbf{f}(\mathbf{A})\|, \quad (5)$$

where $c_1 = O(\|\boldsymbol{\mu}\|)$, $c_2 = O(\mathbb{E}\{\|\mathbf{N}\|\}/\delta)$.

Proof. By substituting $\mathbf{X} = \mathbf{F} + \mathbf{N}$, we obtain $\|\mathbf{f}(\mathbf{A})\mathbf{X} - \mathbf{F}\| \leq \|\mathbf{f}(\mathbf{A})\mathbf{F} - \mathbf{F}\| + \|\mathbf{f}(\mathbf{A})\mathbf{N}\|$. For the first term, by substituting $\mathbf{F} = \mathbf{L}\boldsymbol{\mu}$, we obtain $\|\mathbf{f}(\mathbf{A})\mathbf{F} - \mathbf{F}\| \leq \|\mathbf{f}(\mathbf{A})\mathbf{L} - \mathbf{L}\| O(\|\boldsymbol{\mu}\|)$. For the second term, by Markov inequality, we obtain, $\forall t > 0$, $\text{Prob}\{\|\mathbf{f}(\mathbf{A})\mathbf{N}\| > t\} \leq \frac{\mathbb{E}(\|\mathbf{f}(\mathbf{A})\mathbf{N}\|)}{t}$. By substituting $t = \frac{\mathbb{E}(\|\mathbf{f}(\mathbf{A})\mathbf{N}\|)}{\delta}$, we obtain $\text{Prob}\left\{\|\mathbf{f}(\mathbf{A})\mathbf{N}\| \leq O\left(\frac{\mathbb{E}\{\|\mathbf{N}\|\}}{\delta}\right) \|\mathbf{f}(\mathbf{A})\|\right\} \leq 1 - \delta$.

Theorem 1 suggests that the effect of a graph filter is upper bounded by two parts, as shown on the right side of Equation 5. For the first part, $\|\mathbf{f}(\mathbf{A})\mathbf{L} - \mathbf{L}\|$ represents the bias introduced by the filter $\mathbf{f}(\mathbf{A})$ on the ideal feature matrix $\mathbf{F} = \mathbf{L}\boldsymbol{\mu}$, which is related to the label matrix $\mathbf{L}$, weighted by a constant of the embedding vectors of the labels. For the second part, $\|\mathbf{f}(\mathbf{A})\|$ weighted by the constant related to the noise, reflects the effect that $\mathbf{f}(\mathbf{A})$ filters the noise $\mathbf{N}$. Intuitively, better performance can be achieved when bias is minimized and more noise is filtered out. Our definition of $\mathcal{C}(\mathbf{A})$, which captures the compatibility between graph topology and tasks, aligns with Theorem 1, as it considers both parts. On the one hand, $\mathcal{I}(\mathbf{A})$ corresponds to $\|\mathbf{f}(\mathbf{A})\mathbf{L} - \mathbf{L}\|$, as a smaller value of $\|\mathbf{f}(\mathbf{A})\mathbf{L} - \mathbf{L}\|$ indicates a higher similarity between $\mathbf{f}(\mathbf{A})\mathbf{L}$ and $\mathbf{L}$. On the other hand, $\mathcal{R}(\mathbf{A})$ aligns with $\|\mathbf{f}(\mathbf{A})\|$, as a smaller value of $\|\mathbf{f}(\mathbf{A})\|$ implies less regularization on $\mathbf{f}(\mathbf{A})$. Furthermore, our definition of $\mathcal{C}(\mathbf{A})$ offers more flexibility, as the similarity function $\mathcal{S}(\cdot)$ is not limited to Euclidean distance and the regularization function $\mathcal{R}(\cdot)$ is not limited to L2-norm. This flexibility allows for a broader range of choices in defining the compatibility measure based on specific requirements and characteristics of the graph learning task at hand.

We further investigate the bias introduced by $\mathbf{f}(\mathbf{A})$ on $\mathbf{F}$ by analyzing the change in the distance between a node and its farthest inter-class node, which represents the radius of nodes belonging to a different community distributed, centered on the node. A larger radius indicates easier classification. For noise filtering, we measure the change in variance after applying the filter. In the following, we present the results in detail.

Theorem 2. Suppose that $\mathbf{f}(\mathbf{A}) = \text{RowNorm}(\sum_{k=0}^K \gamma_k \hat{\mathbf{A}}^k)$, $\sum_{k=0}^K \gamma_k = 1$ and $\gamma_k \geq 0$, which are low-pass filters [4, 11]. We define the distance between node v_i and the farthest node to v_i that belongs to a different community as $\mathcal{D}(\mathbf{X}, v_i) = \max_{v_j \in \mathcal{V}, c_i \neq c_j} \|\mathbf{X}_{i,:} - \mathbf{X}_{j,:}\|$. We have

- The maximum distance between v_i and the farthest node to v_i that belongs to a different community decreases after applying the graph filter: $\forall v_i \in \mathcal{V}, \mathcal{D}(\mathbf{F}, v_i) \geq \mathcal{D}(\mathbf{f}(\mathbf{A})\mathbf{F}, v_i)$.
- The total variance of the noise decreases: $\|\text{Var}\{\mathbf{N}\}\|_1 \geq \|\text{Var}\{\mathbf{f}(\mathbf{A})\mathbf{N}\}\|_1$, where $\|\mathbf{X}\|_1 = \sum_i \sum_j |X_{i,j}|$.

Proof. For the first result, by definition,

$$\begin{aligned} \mathcal{D}(\mathbf{f}(\mathbf{A})\mathbf{F}, v_i) &= \max_{v_j \in \mathcal{V}, c_i \neq c_j} \left\| \sum_{k \in \mathcal{V}} (\mathbf{f}(\mathbf{A}))_{i,k} \mathbf{F}_{k,:} - \sum_{k \in \mathcal{V}} (\mathbf{f}(\mathbf{A}))_{j,k} \mathbf{F}_{k,:} \right\| \\ &= \max_{v_j \in \mathcal{V}, c_i \neq c_j} \left\| (\mathbf{f}(\mathbf{A}))_{i,:} \mathbf{L}\mu - (\mathbf{f}(\mathbf{A}))_{j,:} \mathbf{L}\mu \right\| = \max_{v_j \in \mathcal{V}, c_i \neq c_j} \left\| \sum_{c=1}^C l_c^{(i,j)} \mu_c \right\|, \end{aligned}$$

where $l_c^{(i,j)} = (\mathbf{f}(\mathbf{A}))_{i,:} \mathbf{L} - (\mathbf{f}(\mathbf{A}))_{j,:} \mathbf{L}$ and $l_c^{(i,j)} \in \mathbb{R}^{1 \times C}$, and $(\mathbf{f}(\mathbf{A}))_{i,:}$, $(\mathbf{f}(\mathbf{A}))_{j,:}$ represents the i^{th} and j^{th} row of $\mathbf{f}(\mathbf{A})$ respectively. And $\exists w_{m,n}^{(i,j)}, \sum_{c=1}^C l_c^{(i,j)} \mu_c = \sum_{m=1}^C \sum_{n=1}^C w_{m,n}^{(i,j)} (\mu_m - \mu_n)$, such that $\forall 1 \leq c, m, n, \leq C, w_{c,c}^{(i,j)} = 0, w_{c,m}^{(i,j)} w_{c,n}^{(i,j)} \geq 0, \min(w_{m,n}^{(i,j)}, w_{n,m}^{(i,j)}) = 0$. This can be seen as the process that the positive elements in $l_c^{(i,j)}$ distribute its value to negative elements in $l_c^{(i,j)}$. Then we can obtain $\sum_{c=1}^C |l_c^{(i,j)}| = 2 \sum_{m=1}^C \sum_{n=1}^C |w_{m,n}^{(i,j)}|$. Due to the fact that all elements in $\mathbf{f}(\mathbf{A})$ are non-negative and each row sum of $\mathbf{f}(\mathbf{A})$ equals one, we obtain $\forall v_i, v_j \in \mathcal{V}, \sum_{c=1}^C l_c^{(i,j)} = 0$ and $\sum_{c=1}^C |l_c^{(i,j)}| \leq 2$. Further, we denote $v_{j*} = \arg\max_{v_j \in \mathcal{V}, c_i \neq c_j} \|\mathbf{F}_{i,:} - \mathbf{F}_{j,:}\|$ for $\forall v_i \in \mathcal{V}$, then we obtain $\mathcal{D}(\mathbf{F}, v_i) = \|\mathbf{F}_{i,:} - \mathbf{F}_{j*,:}\| = \|\mu_{c_i} - \mu_{c_{j*}}\|$. Subsequently, we derive

$$\begin{aligned} \mathcal{D}(\mathbf{f}(\mathbf{A})\mathbf{F}, v_i) &= \max_{v_j \in \mathcal{V}, c_i \neq c_j} \left\| \sum_{m=1}^C \sum_{n=1}^C w_{m,n}^{(i,j)} (\mu_m - \mu_n) \right\| \\ &\leq \max_{v_j \in \mathcal{V}, c_i \neq c_j} \sum_{m=1}^C \sum_{n=1}^C |w_{m,n}^{(i,j)}| \|(\mu_m - \mu_n)\| \leq \max_{v_j \in \mathcal{V}, c_i \neq c_j} \sum_{m=1}^C \sum_{n=1}^C |w_{m,n}^{(i,j)}| \|\mu_{c_i} - \mu_{c_{j*}}\| \\ &\leq \|\mu_{c_i} - \mu_{c_{j*}}\| = \mathcal{D}(\mathbf{F}, v_i). \end{aligned}$$

For the second result, $\mathbf{f}(\mathbf{A})$ is a row stochastic matrix, so absolute values of all eigenvalues of $\mathbf{f}(\mathbf{A})$ are within 1. Therefore, $\|\text{Var}\{\mathbf{f}(\mathbf{A})\mathbf{N}\}\|_1 = \text{trace}(\mathbb{E}(\mathbf{f}(\mathbf{A})\mathbf{N}\mathbf{N}^\top \mathbf{f}(\mathbf{A})^\top)) \leq n\sigma^2 \cdot \text{trace}(\mathbb{E}(\mathbf{f}(\mathbf{A})\mathbf{f}(\mathbf{A})^\top)) \leq n^2\sigma^2 = \|\text{Var}\{\mathbf{N}\}\|_1$.




Fig. 1: The performance of GCN and MLP on the synthetic graph datasets generated by cSBM with different σ, p, q , corresponding to the noise variance of features, intra-community and inter-community connection probability. The synthetic datasets share the same statistics with Cora [20], such as the number of nodes and edges and the dimension of features. p and q used to generate SBMs are (0.9, 0.1), (0.8, 0.2), and (0.7, 0.3) respectively from left to right.

Theorem 2 shows that the distances analyzed decrease after applying the graph filter, indicating that the difficulty increases for classification and reflects that bias is introduced. For noise, the variance of the noise decreases through the low-pass filter, suggesting that the noise is reduced. In order to analyze the compatibility comprehensively, we need to consider these two effects simultaneously. This also explains why $\mathcal{I}(\mathbf{A})$ and $\mathcal{R}(\mathbf{A})$ are necessary in $\mathcal{C}(\mathbf{A})$.

4.2 Motivating Example

We further present a case study on cSBM to analyze the performance of GCN with different feature noise and show the effects of $\mathbf{f}(\mathbf{A})$.

As shown in Figure 1, MLP outperforms GCN when the noise variance σ is close to zero. The performance gap occurs because $\mathbf{f}(\mathbf{A})$ biases the prediction of GCN and the lower the quality of the topology, i.e., the lower p , the larger the gap, from left to right. However, as σ increases, the performance of both MLP and GCN decreases, but the performance of GCN decreases more slowly than that of MLP because GCN smooths the noise through topology. During this phase, the performance of MLP gradually approaches that of GCN until GCN outperforms MLP. When the noise variance is too large to dominate the feature, the performance of MLP and GCN drops and converges. The convergence performance of GCN is usually higher than MLP, as MLP degenerates into random guesses, while GCN is still able to learn from topology.

This motivating experiment shows that, compared with MLPs which do not utilize graph topology, the bias introduced by the low-pass filter $\mathbf{f}(\mathbf{A})$ results in that MLPs outperform GNNs when the noise variance in features is near zero, suggesting that this effect is generally bad for performance. However, compared to MLPs, low-pass filters $\mathbf{f}(\mathbf{A})$ also provide an additional noise reduction ability, which improves the performance of GNNs when the noise variance in features is greater than zero. Therefore, to fully analyze the compatibility between topology and task in GNN models, we design both $\mathcal{I}(\mathbf{A})$ and $\mathcal{R}(\mathbf{A})$ in $\mathcal{C}(\mathbf{A})$ to analyze and quantify these two effects simultaneously.

5 Experiments

We conduct experiments to validate the performance of TopoInf:

- validate the effectiveness of TopoInf on graphs with ground truth labels;
- estimate TopoInf based on pseudo labels and utilize the estimated TopoInf to refine graph topology for improving model performance;
- show further applications of model improvement, by modifying topology guided by TopoInf, with DropEdge [19] as an example.

Setup. We choose six real-world graph datasets, namely Cora, CiteSeer, PubMed, Computers, Photos and Actor in our experiments [20, 28, 21, 17]. For Computers and Photos, we randomly select 20 nodes from each class for training, 30 nodes from each class for validation, and use the remaining nodes for testing. For other datasets, we use their public train/val/test splits. We choose nine widely adopted models, namely GCN, SGC, APPNP, GAT, GIN, TAGCN, GPRGNN, BernNet and GCNII to work on [12, 26, 13, 25, 27, 10, 4, 11, 3].

5.1 Validating TopoInf on Graphs

Can TopoInf reflect edges’ influence on model performance? We first conduct experiments on TopoInf computed using ground truth labels. Based on the former analysis, the sign of TopoInf implicates the directional influence of removing e_{ij} on the model performance, while the absolute value of TopoInf reflects the magnitude of the influence. We validate these implications in real datasets. For each dataset, we compute TopoInf based on ground truth labels and partition edges with positive TopoInf into the positive set and edges with negative TopoInf into the negative set. Then we sequentially remove edges from the positive/negative set ordered by the absolute values of their TopoInf in a descendant manner and record the performance of models.

As is shown in Figure 2, performance curves exhibit an S-shaped pattern. Specifically, after removing edges in the positive set, the model performance increases, while removing edges in the negative set diminishes the performance. Moreover, removing edges with higher absolute TopoInf values results in a greater derivative of the performance plots. These are consistent with our previous analysis and show the effectiveness of TopoInf.

In addition, we compare our edge removal strategies with other methods. We choose two baseline methods, namely Random and AdaEdge. For Random, we randomly remove edges with equal probability. AdaEdge [2] divides edges between nodes with the same labels into the negative set and different ones into the positive set. AdaEdge randomly removes edges in each set, while TopoInf removes edges in each set in descending order, sorted by the absolute values of their TopoInf.

Figure 3 shows our results. When an equal number of positive/negative edges are removed, the model performance of Random remains unchanged almost, while the model performance of AdaEdge and TopoInf increases/decreases distinctly, indicating a correlation between both the homophily metric and our




Fig. 2: **Performance change while deleting edges by TopoInf.** Horizontal axis’s meaning: zero point corresponds to the original graph without deleting any edges. The absolute value of the coordinate denotes the ratio of deleted edges to all edges. The negative/positive coordinate corresponds to the case that the graph is obtained by deleting edges with negative/positive TopoInf values of the corresponding proportion, in the descendant order of the absolute value of TopoInf. The vertical axis is the accuracy of the node classification in the GNN models. The TopoInf values are calculated with $\mathcal{V}_t$ in Equation 3 set as the test set. For each step, we remove 10% of edges in the negative/positive set.

compatibility metric with the model performance. Moreover, the model performance of TopoInf increases/decreases more significantly than that of AdaEdge. This shows the effectiveness of TopoInf as it not only identifies the direction of influence but also measures the magnitude of influence.

5.2 Validating Estimated TopoInf

When ground truth labels are generally only available on training sets, we employ an initial training phase using the GNN model to obtain pseudo labels for each node. Thereafter, we estimate TopoInf using these pseudo labels for the nodes without true labels and remove edges based on the estimated TopoInf.

How well does estimated TopoInf work to improve model performance? We conduct experiments on nine GNN models. We train the model and obtain pseudo labels. After that, we can estimate TopoInf based on pseudo labels and refine the original topology based on the estimated TopoInf. Specifically, we remove a given number of edges, which is a hyper-parameter determined by validation, with the highest estimated TopoInf. Since we have already obtained the model parameters through the initial training, we can choose to either continue to apply these parameters on the refined topology (denoted as w/retrain) or retrain the model on the refined topology (denoted as w/o retrain).

Table 2 shows the results. We can see that in almost all cases, the refined topologies based on estimated TopoInf outperform the original ones, in




Fig. 3: **GCN performance with different edge deletion strategies on Cora.** The left figure shows the results of removing positive edges to increase performance (higher is better), and the right figure shows the results of removing negative edges to decrease performance (lower is better). The horizontal axis denotes the ratio of deleted edges to all edges. The vertical axis is the accuracy of the node classification in GCN. The TopoInf values are calculated with $\mathcal{V}_t$ in the Equation 3 set as a test set.

both retraining and non-retraining settings. On SOTA models, such as BernNet, TAGCN and GCNII, TopoInf still achieves competitive results or improves performance. This experiment shows that the estimated TopoInf is still an effective metric for enhancing topology.

Are there any other approaches to utilize estimated TopoInf? In this work, we also demonstrate that the estimated TopoInf can be combined with other methods to guide topology modification and improve model performance. We take DropEdge [19] as an example, where edges are randomly removed from the original graph at each training epoch to help prevent over-fitting and over-smoothing and achieve better performance. We replace the random dropping edge scheme with a TopoInf-guided scheme, where edges with higher TopoInf are more likely to be dropped. Specifically, for TopoInf-guided DropEdge, we define our edge-dropping probability as $P_{ij} \propto \mathbf{A}_{ij} \cdot \exp(\nabla \mathcal{C}_{\mathbf{A}}(e_{ij})/\tau)$, where $\mathbf{A}_{ij}$ indicates the existence of e_{ij} , $\nabla \mathcal{C}_{\mathbf{A}}(e_{ij})$ is the TopoInf of e_{ij} , and τ is the temperature, which is a hyper-parameter and determined by validation.

Table 3 shows the results. In almost all cases, our TopoInf-guided DropEdge method achieves the highest accuracy, and in all cases, our method outperforms DropEdge. Even in cases where DropEdge degenerates the performance of GNNs, such as BernNet on CiteSeer, and GPRGNN on PubMed, TopoInf-guided DropEdge can still improve the performance. This experiment shows the potential to combine TopoInf with other methods for topology modification and performance improvement.

Table 2: **Performance after deleting edges based on estimated TopoInf on different models and datasets.** The estimated TopoInf values are calculated with V_i in Equation 3 set as the whole nodes. The ratio of the edges removed to all edges is a hyperparameter chosen from $\{0.02, 0.04, 0.06, 0.08\}$ and is determined by validation. We run experiments 10 times and report testing accuracy with a 95% confidence interval.

Datasets	Method	GCN	SGC	APPNP	GAT	GIN	TAGCN	GPRGNN	BERNNET	GCNII
CORA	original	80.3 $\pm$ 0.3	80.8 $\pm$ 0.4	80.4 $\pm$ 0.4	80.4 $\pm$ 0.6	74.8 $\pm$ 0.8	82.0$\pm$0.3	81.3 $\pm$ 0.7	82.2 $\pm$ 0.3	83.6$\pm$0.4
	w/o retrain(our)	81.6$\pm$0.4	81.2$\pm$0.5	81.1$\pm$0.5	80.9 $\pm$ 0.8	77.1$\pm$0.8	82.0$\pm$0.2	81.6$\pm$0.6	82.6$\pm$0.4	83.6$\pm$0.1
	w retrain(our)	81.6$\pm$0.4	81.2$\pm$0.4	81.0$\pm$0.4	81.5$\pm$0.4	77.0$\pm$0.8	82.0$\pm$0.2	81.6$\pm$0.6	82.6$\pm$0.4	83.5$\pm$0.1
	original	68.7 $\pm$ 0.5	65.6 $\pm$ 1.0	65.1 $\pm$ 0.5	68.1 $\pm$ 0.7	61.6 $\pm$ 1.4	69.8 $\pm$ 0.4	69.4 $\pm$ 2.1	71.2$\pm$0.3	72.4 $\pm$ 0.6
CITSEER	w/o retrain(our)	70.0$\pm$0.7	66.8 $\pm$ 1.0	67.3 $\pm$ 0.8	70.5$\pm$1.1	62.1$\pm$1.4	70.6 $\pm$ 0.3	71.3$\pm$0.6	71.0 $\pm$ 0.7	72.7$\pm$0.6
	w retrain(our)	69.8 $\pm$ 0.7	66.9$\pm$1.2	67.4$\pm$0.8	70.5$\pm$1.0	62.1$\pm$1.4	70.8$\pm$0.2	71.3$\pm$0.5	71.0 $\pm$ 0.6	72.7$\pm$0.5
	original	77.9 $\pm$ 0.3	77.6 $\pm$ 0.2	77.6 $\pm$ 0.4	77.4 $\pm$ 0.3	75.5 $\pm$ 0.6	79.4 $\pm$ 0.1	78.6 $\pm$ 0.5	78.9 $\pm$ 0.3	79.2 $\pm$ 0.2
	w/o retrain(our)	78.7 $\pm$ 0.3	78.5$\pm$0.5	78.5$\pm$0.4	77.8$\pm$0.2	76.0 $\pm$ 0.6	80.0$\pm$0.3	79.3$\pm$0.5	79.2$\pm$0.2	79.4$\pm$0.2
PUBMED	w retrain(our)	78.8$\pm$0.2	78.4 $\pm$ 0.5	78.5$\pm$0.3	77.9$\pm$0.2	76.1$\pm$0.7	79.7 $\pm$ 0.3	79.2 $\pm$ 0.5	79.1 $\pm$ 0.5	79.3 $\pm$ 0.2

Table 3: **Performance of different DropEdge strategies on different datasets and models.** The estimated TopoInf values are calculated with V_i in Equation 3 set as the whole nodes. The DropEdge rate and the temperature are hyper-parameters chosen from $\{0.3, 0.4, 0.5\}$ and $\{0.5, 0.75, 1\}$ respectively and are determined by validation. We run experiments 10 times and report testing accuracy with a 95% confidence interval.

Datasets	DropEdge	GCN	SGC	APPNP	GAT	GIN	TAGCN	GPRGNN	BERNNET	GCNII
CORA	without	80.3 $\pm$ 0.4	80.7 $\pm$ 0.7	80.5 $\pm$ 0.5	80.6 $\pm$ 0.6	74.7 $\pm$ 0.8	82.0 $\pm$ 0.4	81.5 $\pm$ 1.7	82.2 $\pm$ 0.4	83.6 $\pm$ 0.5
	random	81.9 $\pm$ 0.7	81.4 $\pm$ 0.8	80.8 $\pm$ 0.7	81.7 $\pm$ 0.6	76.1 $\pm$ 0.9	83.0$\pm$0.4	82.0 $\pm$ 0.9	82.3 $\pm$ 0.2	84.2 $\pm$ 0.4
	TopoInf(our)	82.2$\pm$0.6	81.5$\pm$0.9	81.6$\pm$0.5	82.2$\pm$0.5	77.2$\pm$0.7	83.0$\pm$0.3	82.1$\pm$0.8	82.8$\pm$0.3	84.5$\pm$0.4
	without	68.8 $\pm$ 0.7	65.9 $\pm$ 1.3	65.1 $\pm$ 0.6	69.3 $\pm$ 0.9	61.4 $\pm$ 1.4	69.9 $\pm$ 0.4	70.4 $\pm$ 0.7	71.2 $\pm$ 0.5	72.5 $\pm$ 0.6
CITSEER	random	68.7 $\pm$ 0.9	67.1 $\pm$ 1.0	67.8 $\pm$ 0.8	69.8 $\pm$ 0.8	62.4 $\pm$ 1.0	71.2 $\pm$ 0.4	70.9 $\pm$ 0.7	71.1 $\pm$ 0.6	72.5 $\pm$ 0.6
	TopoInf(our)	69.3$\pm$1.1	67.8$\pm$0.9	68.3$\pm$0.9	70.3$\pm$0.5	63.5$\pm$0.9	71.7$\pm$0.3	71.1$\pm$0.6	71.8$\pm$0.4	72.9$\pm$0.6
	without	77.9 $\pm$ 0.5	77.6 $\pm$ 0.2	77.6 $\pm$ 0.4	77.5$\pm$0.5	75.5 $\pm$ 0.7	79.5 $\pm$ 0.3	78.7 $\pm$ 0.7	78.9$\pm$0.4	79.2$\pm$0.2
	random	77.5 $\pm$ 0.4	77.9 $\pm$ 0.3	77.6 $\pm$ 0.5	77.0 $\pm$ 0.4	76.9 $\pm$ 0.5	79.2 $\pm$ 0.3	78.5 $\pm$ 0.7	78.7 $\pm$ 0.5	78.8 $\pm$ 0.3
PUBMED	TopoInf(our)	78.0$\pm$0.6	78.3$\pm$0.3	78.0$\pm$0.4	77.3$\pm$0.4	77.4$\pm$0.7	79.7$\pm$0.2	79.3$\pm$0.4	78.8 $\pm$ 0.5	79.1 $\pm$ 0.2

6 Related Works

Graph filters optimization. These studies approach graphs as filters for features and aim to find the optimal filter coefficients for the graph learning task. ChebNet [6] approximates the spectral graph convolutions using Chebyshev polynomials. GPRGNN [4] adaptively learns a polynomial filter and jointly optimizes

the neural parameters of the neural network as well as filter coefficients. ASGC [1] obtains filter coefficients by minimizing the approximation error between the raw feature and the feature filtered by graph topology, to make SGC appropriate in heterophilic settings. These works improve graph learning by optimizing coefficients of graph filters, while we focus on optimizing graph topology.

Graph topology optimization. Curvature-based methods [23] view graphs from a geometric view, measure the effect of edges by the graph curvature, and rewire the graph topology with the help of curvature, to improve the model performance. Node pair distances [2, 24] and edge signal-to-noise ratio [9] are proposed to measure the topological information and alleviate the over-smoothing to further build deeper models to exploit the multi-hop neighborhood structures. In the era of large models, Sun et al. [22] utilize LLMs to evaluate the edges between nodes based on node attributes in order to optimize graph topology.

7 Conclusion

In this work, we model and analyze the compatibility between the topological information of graph data and downstream task objectives, and propose a new metric called TopoInf to characterize the influence of one single edge by measuring the change of the compatibility. We also discuss the potential applications of leveraging TopoInf to model performance adjustment and empirically demonstrate the effectiveness of TopoInf. In future, a better estimation of the topological effect of GNNs may further improve our metric.

Acknowledgment

This work was supported by the National Key Research and Development Plan No. 2022YFB3904204, NSF China under Grant No. 62202299, 62020106005, 61960206002, Shanghai Natural Science Foundation No. 22ZR1429100.

References

1. Chanpuriya, S., Musco, C.: Simplified graph convolution with heterophily. In: NeurIPS (2022)
2. Chen, D., Lin, Y., Li, W., et al.: Measuring and relieving the over-smoothing problem for graph neural networks from the topological view. In: AAAI (2020)
3. Chen, M., Wei, Z., Huang, Z., et al.: Simple and deep graph convolutional networks. In: ICML (2020)
4. Chien, E., Peng, J., Li, P., Milenkovic, O.: Adaptive universal generalized pagerank graph neural network. In: ICLR (2021)
5. Chin, A., Chen, Y., M. Altenburger, K., Ugander, J.: Decoupled smoothing on graphs. In: WWW (2019)
6. Defferrard, M., Bresson, X., Vandergheynst, P.: Convolutional neural networks on graphs with fast localized spectral filtering. In: NeurIPS (2016)

7. Deshpande, Y., Montanari, A., Mossel, E., Sen, S.: Contextual stochastic block models. In: NeurIPS (2018)
8. Dong, H., Chen, J., Feng, F., et al.: On the equivalence of decoupled graph convolution network and label propagation. In: WWW (2021)
9. Dong, M., Kluger, Y.: Towards understanding and reducing graph structural noise for GNNs. In: ICML (2023)
10. Du, J., Zhang, S., Wu, G., Moura, J.M.F., Kar, S.: Topology adaptive graph convolutional networks. arXiv preprint arXiv:1710.10370 (2018)
11. He, M., Wei, Z., Huang, Z., Xu, H.: Bernnet: Learning arbitrary graph spectral filters via Bernstein approximation. In: NeurIPS (2021)
12. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. In: ICLR (2017)
13. Klicpera, J., Bojchevski, A., Günnemann, S.: Predict then propagate: Graph neural networks meet personalized PageRank. In: ICLR (2019)
14. Luo, D., Cheng, W., Yu, W., et al.: Learning to drop: Robust graph neural network via topological denoising. In: WSDM (2021)
15. Ma, Y., Liu, X., Shah, N., et al.: Is homophily a necessity for graph neural networks? arXiv preprint arXiv:2106.06134 (2021)
16. Nt, H., Maehara, T.: Revisiting graph neural networks: All we have is low-pass filters. arXiv preprint arXiv:1905.09550 (2019)
17. Pei, H., Wei, B., Chang, K.C., Lei, Y., Yang, B.: Geom-gcn: Geometric graph convolutional networks. In: ICLR (2020)
18. Ren, Y., Bai, J., Zhang, J.: Label contrastive coding based graph neural network for graph classification. In: DASFAA (2021)
19. Rong, Y., Huang, W., Xu, T., Huang, J.: Dropedge: Towards deep graph convolutional networks on node classification. In: ICLR (2020)
20. Sen, P., Namata, G., Bilgic, M., et al.: Collective classification in network data. *AI magazine* (2008)
21. Shchur, O., Mumme, M., Bojchevski, A., Günnemann, S.: Pitfalls of graph neural network evaluation. arXiv preprint arXiv:1811.05868 (2018)
22. Sun, S., Yuxiang, R., et al.: Large language models as topological structure enhancers for text-attributed graphs. arXiv preprint arXiv:2311.14324 (2024)
23. Topping, J., Di Giovanni, F., Chamberlain, B.P., et al.: Understanding over-squashing and bottlenecks on graphs via curvature. In: ICLR (2022)
24. Bai, J., Ren, Y., Zhang, J.: Measuring and sampling: A metric-guided subgraph learning framework for graph neural network. *Int. J. Intell. Syst.* (2022)
25. Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., Bengio, Y.: Graph attention networks. In: ICLR (2018)
26. Wu, F., Souza, A., Zhang, T., Fifty, C., Yu, T., Weinberger, K.: Simplifying graph convolutional networks. In: ICML (2019)
27. Xu, K., Hu, W., Leskovec, J., Jegelka, S.: How powerful are graph neural networks? In: ICLR (2019)
28. Yang, Z., Cohen, W., Salakhudinov, R.: Revisiting semi-supervised learning with graph embeddings. In: WWW (2016)
29. Zhang, W., Yang, M., Sheng, Z., et al.: Node-dependent local smoothing for scalable graph learning. In: NeurIPS (2021)
30. Zhu, H., Koniusz, P.: Simple spectral graph convolution. In: ICLR (2021)
31. Zhu, J., Rossi, R.A., Rao, A., et al.: Graph neural networks with heterophily. In: AAAI (2021)
32. Zhu, J., Yan, Y., et al.: Beyond homophily in graph neural networks: Current limitations and effective designs. In: NeurIPS (2020)