

Confidence May Cheat: Self-Training on Graph Neural Networks under Distribution Shift

Hongrui Liu*
liuhongrui@bupt.edu.cn
Beijing University of Posts and
Telecommunications

Binbin Hu
bin.hbb@antfin.com
Ant Group

Xiao Wang
xiaowang@bupt.edu.cn
Beijing University of Posts and
Telecommunications
Peng Cheng Laboratory

Chuan Shi†
shichuan@bupt.edu.cn
Beijing University of Posts and
Telecommunications
Peng Cheng Laboratory

Zhiqiang Zhang
lingyao.zzq@antfin.com
Ant Group

Jun Zhou
jun.zhoujun@antfin.com
Ant Group

ABSTRACT

Graph Convolutional Networks (GCNs) have recently attracted vast interest and achieved state-of-the-art performance on graphs, but its success could typically hinge on careful training with amounts of expensive and time-consuming labeled data. To alleviate labeled data scarcity, self-training methods have been widely adopted on graphs by labeling high-confidence unlabeled nodes and then adding them to the training step. In this line, we empirically make a thorough study for current self-training methods on graphs. Surprisingly, we find that high-confidence unlabeled nodes are not always useful, and even introduce the distribution shift issue between the original labeled dataset and the augmented dataset by self-training, severely hindering the capability of self-training on graphs. To this end, in this paper, we propose a novel **Distribution Recovered Graph Self-Training** framework (DR-GST), which could recover the distribution of the original labeled dataset. Specifically, we first prove the equality of loss function in self-training framework under the distribution shift case and the population distribution if each pseudo-labeled node is weighted by a proper coefficient. Considering the intractability of the coefficient, we then propose to replace the coefficient with the information gain after observing the same changing trend between them, where information gain is respectively estimated via both dropout variational inference and droppedge variational inference in DR-GST. However, such a weighted loss function will enlarge the impact of incorrect pseudo labels. As a result, we apply the loss correction method to improve the quality of pseudo labels. Both our theoretical analysis and extensive experiments on five benchmark datasets demonstrate the effectiveness of the proposed DR-GST, as well as each well-designed component in DR-GST.

*Work done during internship at Ant Group.

†Corresponding author

CCS CONCEPTS

• **Computing methodologies** → **Neural networks**; • **Theory of computation** → **Social networks**; **Semi-supervised learning**.

KEYWORDS

Graph Neural Networks, Self-Training, Information Gain

ACM Reference Format:

Hongrui Liu, Binbin Hu, Xiao Wang, Chuan Shi, Zhiqiang Zhang, and Jun Zhou. 2022. Confidence May Cheat: Self-Training on Graph Neural Networks under Distribution Shift. In *Proceedings of the ACM Web Conference 2022 (WWW '22)*, April 25–29, 2022, Virtual Event, Lyon, France. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3485447.3512172>

1 INTRODUCTION

Graphs are ubiquitous across many real-world applications, ranging from citation and social network analysis to protein interface and chemical bond prediction. With the surge of demands, Graph Convolution Network (GCN) and its variants [17, 18, 30, 32, 35] (abbreviated as GCNs) have recently attracted vast interest and achieved state-of-the-art performance in various tasks on graphs, most notably semi-supervised node classification. Nevertheless, its success could typically hinge on careful training with large amounts of labeled data, which is expensive and time-consuming to be obtained [28]. Empirically, the performance of GCNs will rapidly decline with the decrease of labeled data [38].

As one of the promising approaches, self-training [16, 19] aims at addressing labeled data scarcity by making full use of abundant unlabeled data in addition to task-specific labeled data. Given an arbitrary model trained on the original labeled data as the *teacher* model, the key idea of self-training is to pseudo-label high-confidence unlabeled samples to augment the above labeled data, and a *student* model is trained with augmented data to replace the *teacher* model. Such an iteration learning is repeated until convergence¹. Analogously, self-training has great potential to facilitate advancing GCNs to exploiting unlabeled data [20, 28, 38]. Whereas, these studies only focus on the high-confidence nodes on account of the prefabricated assumption that the higher the confidence, the more

¹The *teacher-student* term is commonly adopted in current self-training studies [8, 14, 22], and we just reuse it here for a clearer explanation.

accurate the prediction. Naturally, we are curious about such a fundamental question, “Are all the unlabeled nodes pseudo-labeled with high confidence truly useful?”

As a motivating example, we conduct an analysis experiment on a benchmark dataset Cora [26] to explore how much additional information these high-confidence nodes can bring to the model (denoted as information gain). More details can be seen in Section 3. Surprisingly, our experimental results show a clear negative correlation between the confidence and the information gain, implying that nodes pseudo-labeled by existing graph self-training methods with high confidence may be low-information-gain and useless. To further understand the underlying reason, we illustrate the distribution of unlabeled nodes and find these high-confidence (or low-information-gain) nodes are far from the decision boundary, which implies that they potentially guide the model to perform worthless optimization for a more crisp decision boundary. Existing graph self-training methods which focus on high-confidence nodes are “cheated” by confidence in this way.

In light of the above observations, we further investigate into what will happen when self-training is cheated by confidence. We discover that during the optimization procedure dominated by *easy nodes* (i.e., nodes with low information gain), the *Distribution Shift* phenomenon between the original and augmented dataset gradually appears. This is because more and more easy nodes selected by high confidence are added to the original labeled dataset, leading to the distribution gradually shifting to the augmented dataset and overmuch attention paid on such easy nodes as a result. Not surprisingly, this issue will severely threaten the capacity of self-training on graphs, since the distribution of the augmented dataset is different from the population distribution, resulting in a terrible generalization during evaluation. Alleviating distribution shift from self-training on graphs is in urgent demand, which is unexplored in existing studies.

In this paper, we propose an information gain weighted self-training framework DR-GST which could recover the distribution of original labeled dataset. Specifically, we first prove that the loss function of the self-training framework under the distribution shift case is equal to that under the population distribution if we could weight each pseudo-labeled node with a proper coefficient. But the coefficient is generally intractable in practice. Then we discover the same changing trend between the coefficient and information gain, and propose to replace the coefficient with information gain, where the information gain can be estimated via both dropout variational inference and droppedge variational inference. Consequently, we can recover the shifted distribution with the newly proposed information gain weighted loss function. Such a loss function forces the model to pay more attention to *hard nodes*, i.e., nodes with high information gain, but will enlarge the impact of incorrect pseudo labels. Therefore, we apply loss correction [10, 23, 27] to self-training to correct the prediction of the student model, so that the impact of incorrect pseudo labels from the teacher model can be alleviated in this way. Finally, we conduct a theoretical analysis of self-training on graphs, and the conclusion shows both distribution shift and incorrect pseudo labels will severely hinder its capability, which is consistent with our designs.

In summary, the main contributions are highlighted as follows:

- We make a thorough study on graph self-training, and find two phenomena below: 1) pseudo-labeled high-confidence nodes may cheat. 2) distribution shift between the original labeled dataset and the augmented dataset. Both of them severely hinder the capability of self-training on graphs.
- We propose a novel graph self-training framework DR-GST that not only addresses the distribution shift issue from the view of information gain, but also is equipped with the creative loss correction strategy for improving qualities of pseudo labels.
- We theoretically analyze the rationality of the whole DR-GST framework and extensive experimental results on five benchmark datasets demonstrates that DR-GST consistently and significantly outperforms various state-of-arts.

2 PRELIMINARY

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{X})$ be a graph with the adjacent matrix $\mathbf{A} \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{V}|}$, where \mathcal{V} and \mathcal{E} are respectively the set of nodes and edges, and $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{|\mathcal{V}|}] \in \mathbb{R}^{|\mathcal{V}| \times D_v}$ is the D_v -dimensional feature matrix for nodes. In the common semi-supervised node classification setting, we only have access to a small amounts of labeled nodes \mathcal{V}_L with their labels \mathcal{Y}_L along with a larger amounts of unlabeled nodes \mathcal{V}_U , where $|\mathcal{V}_L| \ll |\mathcal{V}_U|$.

Self-training Generally, self-training methods on graphs firstly train a vanilla GCN as the base *teacher* model $f_\theta(\mathbf{X}, \mathbf{A})$ with ground-truth labels \mathcal{Y}_L , where θ is the model parameter set. We could obtain the probability vector for each node $v_i \in \mathcal{V}$ as: $\mathbf{p}(y_i | \mathbf{x}_i, \mathbf{A}; \theta) = f_\theta(\mathbf{x}_i, \mathbf{A})$. For convenience, we abbreviate it to \mathbf{p}_i and denote the j -th element of \mathbf{p}_i by $p_{i,j}$. Next, the teacher model pseudo-labels a subset $\mathcal{S}_U \subset \mathcal{V}_U$ of unlabeled nodes with its prediction $\tilde{y}_u = \arg \max_j p_{u,j}$ for each node $v_u \in \mathcal{S}_U$. The selection of \mathcal{S}_U is based on the confidence score $r_i = \max_j p_{i,j}$, i.e., only nodes with r_i higher than a threshold or top- k high-confidence nodes are added to the labeled dataset. Then the augmented dataset $\mathcal{V}_L \cup \mathcal{S}_U$ is used to train a *student* model $f_{\tilde{\theta}}$ with the following objective function.

$$\begin{aligned} \min_{\theta \in \Theta} \mathcal{L}(\mathbf{A}, \mathbf{X}, \mathcal{Y}_L) = & \min_{\theta \in \Theta} \mathbb{E}_{v_i \in \mathcal{V}_L, y_i \in \mathcal{Y}_L} l(y_i, \mathbf{p}_i) \\ & + \lambda \mathbb{E}_{v_u \in \mathcal{S}_U, \mathcal{S}_U \subset \mathcal{V}_U} \mathbb{E}_{\tilde{y}_u \sim \mathbf{p}(y_u | \mathbf{x}_u, \mathbf{A}; \theta)} l(\tilde{y}_u, \mathbf{p}_u), \end{aligned} \quad (1)$$

where $l(y_i, \mathbf{p}_i) = -\log p_{i,y_i}$ is the multi-class cross entropy loss and we fix $\lambda = 1$ in this paper. Finally we replace the teacher model with the student model and iterate the above procedure until convergence.

Information Gain As can be seen in Eq. 1, self-training on graphs will exploit the unlabeled data to train the whole model. Here, we aim to measure how an unlabeled node contributes to the model optimization in a principled way, i.e., information gain. Information gain usually measures the reduction in information given a random variable, where information is generally calculated by the Shannon’s entropy [6]. We utilize the information gain here to seek the node v_u which owns the most information about parameters θ of model posterior and could reduce the number of possible parameter hypotheses maximally fast. We refer to this type of information gain as information gain about model parameters [22]. Formally, given a node v_u , the information gain about model

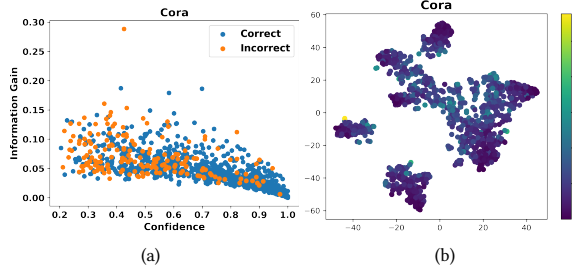


Figure 1: (a): Relationship between confidence and information gain on Cora. (b): Visualization of embeddings on Cora

parameters is defined as \mathbb{B}_u , which could be calculated as follows:

$$\mathbb{B}_u(y_u, \theta | \mathbf{x}_u, \mathbf{A}, \mathcal{G}) = \mathbb{H}[\mathbb{E}_{P(\theta|\mathcal{G})}[y_u | \mathbf{x}_u, \mathbf{A}; \theta]] - \mathbb{E}_{P(\theta|\mathcal{G})}[\mathbb{H}[y_u | \mathbf{x}_u, \mathbf{A}; \theta]], \quad (2)$$

where $\mathbb{H}(\cdot)$ denotes the Shannon’s entropy and $P(\theta|\mathcal{G})$ is the distribution of model posterior. The first term measures the information of the model parameters under posterior, while the second term captures the information of model parameters given an additional node v_u . Obviously, by calculating the difference between the two terms above, \mathbb{B}_u can measure how much information v_u can bring to learn the model parameters θ .

3 EMPIRICAL ANALYSIS

In this section, we conduct a series of empirical analysis to examine whether current graph self-training approaches adopt a principled way to leverage unlabeled data for semi-supervised node classification.

Empirical Analysis of Confidence To better understand the capacity of high-confidence nodes in current self-training approaches, we aim to closely examine that how much additional information these nodes can bring to the model based on information gain. We first visualize the relationship between confidence and information gain in Fig. 1(a), where the x-axis is the confidence while the y-axis is the information gain, and the blue and orange dots respectively represent nodes with correct and incorrect predictions. From Fig. 1(a) we can observe a negative correlation, implying that existing graph self-training methods only focus on *easy* nodes (nodes with low information gain) and confidence may be cheating as a result. Essentially, such a “cheating” phenomenon lies in the worthless optimization for a more crisp decision boundary. Specifically, as shown in Fig. 1(b), on the Cora dataset, we visualize the node embeddings on the last layer of the standard GCN before *softmax* using *t*-SNE [29] algorithm, where a darker dot represents a node with lower information gain. From the plots, we find that most of easy nodes (*i.e.*, low information gain) are far from the decision boundary. Whereas, these nodes are always emphasized by current self-training methods on graphs [20, 28, 38] by force of high confidence. That is, these methods are “cheated” by confidence in this way.

Empirical Analysis of Distribution Shift Furthermore, we investigate what will happen when self-training has been cheated by confidence. As an illustrative example, we randomly generate




Figure 2: Visualization of labeled nodes under the ideal condition. (P_{pop} : distribution before self-training, P_{st} : distribution after self-training)

500 nodes (blue) following two-dimensional Gaussian distribution $\mathcal{N}(0, 0, 0.3, 0.3, 0)$ to represent labeled nodes in one class, and another 4000 nodes (grey) following the distribution of concentric circles [1] to represent labeled nodes belonging to other classes, as shown in Fig. 2(a). Furthermore, following the common self-training setting, a large amount of unlabeled nodes still exists in the dataset, but for clarity, we omit them in the figure. In line with the core idea of current self-training methods, for the “blue” class, unlabeled nodes around the center are pseudo-labeled for self-training since these nodes have high confidence (*a.k.a.*, far from the decision boundary). During iteration, as shown in Fig 2(b), the data distribution will become more and more sharpen since nodes far from the decision boundary are paid disproportionate attention and thus the unsatisfying *Distribution Shift* phenomenon between the original and augmented dataset indeed appears.

4 THE DR-GST FRAMEWORK

In this section, we elaborate the proposed DR-GST, a novel self-training framework aiming at recovering the shifted distribution.

4.1 Information Gain Weighted Loss Function Towards Distribution Shift

We start with the formulation of the self-training task by analyzing the corresponding loss functions. Specifically, assuming that the original labeled dataset follows the population distribution P_{pop} , given a classifier f_θ parameterized by θ , the best parameter set θ could be obtained via minimizing the following loss function:

$$\mathcal{L}_{pop} = \mathbb{E}_{(v_i, y_i) \sim P_{pop}(\mathcal{V}, \mathcal{Y})} l(y_i, \mathbf{p}_i). \quad (3)$$

Similarly, under the distribution shift case caused by self-training, the loss function can be represented as

$$\begin{aligned} \mathcal{L}_{st} = & \frac{|\mathcal{V}_L|}{|\mathcal{V}_L \cup \mathcal{S}_U|} \mathbb{E}_{(v_i, y_i) \sim P_{pop}(\mathcal{V}, \mathcal{Y})} l(y_i, \mathbf{p}_i) \\ & + \frac{|\mathcal{S}_U|}{|\mathcal{V}_L \cup \mathcal{S}_U|} \mathbb{E}_{(v_u, y_u) \sim P_{st}(\mathcal{V}, \mathcal{Y})} l(\bar{y}_u, \mathbf{p}_u), \end{aligned} \quad (4)$$

where P_{st} represents the shifted distribution of the augmented dataset.

Generally, the distribution shift could lead to a terrible generalization during evaluation, and thus severely threaten the capacity of graph self-training. Therefore, It is ideal to optimize f_θ with the loss function \mathcal{L}_{pop} under the population distribution rather than

\mathcal{L}_{st} under the distribution shift case. However, only \mathcal{L}_{st} is available in practice. To close the gap, we show the following theorem.

THEOREM 4.1. *Given \mathcal{L}_{pop} and \mathcal{L}_{st} defined in Eq. 3 and Eq. 4, assuming that $\tilde{y}_u = y_u$ for each pseudo-labeled node $v_u \in \mathcal{S}_U$, then $\mathcal{L}_{st} = \mathcal{L}_{pop}$ holds true if \mathcal{L}_{st} can be written with an additional weight coefficient $\gamma_u = \frac{P_{pop}(v_u, y_u)}{P_{st}(v_u, y_u)}$ as follows:*

$$\mathcal{L}_{st} = \frac{|\mathcal{S}_U|}{|\mathcal{V}_L \cup \mathcal{S}_U|} \mathbb{E}_{(v_u, y_u) \sim P_{st}(\mathcal{V}, \mathcal{Y})} \gamma_u l(\tilde{y}_u, \mathbf{p}_u) + \frac{|\mathcal{V}_L|}{|\mathcal{V}_L \cup \mathcal{S}_U|} \mathbb{E}_{(v_i, y_i) \sim P_{pop}(\mathcal{V}, \mathcal{Y})} l(y_i, \mathbf{p}_i), \quad (5)$$

PROOF. Please refer to A.1.1. \square

Based on Theorem 4.1, we can find that our desired \mathcal{L}_{pop} can be written as the available \mathcal{L}_{st} only if a coefficient γ_u is added to \mathcal{L}_{st} . In other words, the distribution shift issue could be addressed by optimizing f_θ with available \mathcal{L}_{st} weighted by γ_u (in Eq. 5). However, it should be noted that the population distribution P_{pop} in \mathcal{L}_{st} is generally intractable, which means that γ_u cannot be accurately calculated.

To this end, we propose to build the bridge between γ_u and the information gain, which is motivated as follows. Recalling the data distributions shown in Fig. 2(a) and Fig. 2(b), we could formally represent the former as P_{pop} and the latter as P_{st} . We visualize the desired weight coefficient $\gamma_u = \frac{P_{pop}(v_u, y_u)}{P_{st}(v_u, y_u)}$ for each pseudo-labeled node v_u in Fig. 2(c) for better understanding its changing trend, where the darker area means the larger γ_u . Obviously, we observe that γ_u becomes smaller when getting closer to the center area (a.k.a., far away from the decision boundary), which is consistent with the change trend of the information gain. This finding inspires us to adopt the information gain to approximate γ_u .

4.2 Information Gain Estimation on Graphs

Next, we elaborate the estimation of the information gain for each node v_u in graph. As mentioned in Eq. 2, the distribution of model posterior $P(\theta|\mathcal{G})$ is desired for calculating information gain, but it is intractable in practice, and always computationally expensive for traditional bayesian neural networks [2, 11, 15]. Instead, we could shift attention towards dropout [9] and dropedge [25], a type of regularization technique for preventing over-fitting and over-smoothing in GCNs, which could be both interpreted as an approximation of $P(\theta|\mathcal{G})$ [13]. Consequently, we propose to estimate the information gain assisted with dropout and dropedge (a.k.a., *dropout* and *dropedge variational inference*), which takes into account both features and the network topology in our unified framework DR-GST. For distinction, we refer to DR-GST with dropout variational inference as DR-GST_{do} and that with dropedge variational inference as DR-GST_{de}.

4.2.1 Dropout Variational Inference. Specifically, given a L -layer GCN model f_θ , its l -th layer output $\mathbf{H}^{(l)} \in \mathbb{R}^{|\mathcal{V}| \times D_l}$ can be obtained by

$$\mathbf{H}^{(l)} = \sigma(\mathfrak{N}(\mathbf{A})\mathbf{H}^{(l-1)}\mathbf{W}^{(l-1)}), \quad (6)$$

where $\mathfrak{N}(\cdot)$ represents the normalizing operator, $\mathbf{W}^{(l-1)} \in \mathbb{R}^{D_{l-1} \times D_l}$ is the $(l-1)$ -th layer weight matrix, $\sigma(\cdot)$ is the activation function

and $\mathbf{H}^{(1)} = \mathbf{X} \in \mathbb{R}^{|\mathcal{V}| \times D_v}$, $\theta = \{\mathbf{W}^{(l)}\}_{l=1}^L$. Dropout randomly masks features of nodes in the graph through drawing from an independent Bernoulli random variable. Formally, the l -th layer output of f_θ with dropout can be written as:

$$\mathbf{H}^{(l)} = \sigma(\mathfrak{N}(\mathbf{A})(\mathbf{H}^{(l-1)} \odot \mathbf{Z}^{(l-1)})\mathbf{W}^{(l-1)}), \quad (7)$$

where each element of $\mathbf{Z}^{(l)} \in \{0, 1\}^{D_{l-1} \times D_l}$ is a sample of Bernoulli random variable, representing whether or not the corresponding feature in $\mathbf{H}^{(l-1)}$ is set to zero.

Such Bernoulli random sampling on features can also be treated as a sample from $P(\theta|\mathcal{G})$ [9], thus we can perform T -times Monte-Carlo sampling (referred to Monte-Carlo dropout, *MC-dropout*) during *inference* to estimate $P(\theta|\mathcal{G})$. At each time t , a probability vector $\tilde{\mathbf{p}}_u^t = \tilde{\mathbf{p}}^t(y_u | \mathbf{x}_u, \mathbf{A}; \tilde{\theta}_t)$ can be obtained by performing forward pass under such a sample weight $\tilde{\theta}_t$, i.e., $\tilde{\mathbf{p}}_u^t = f_{\tilde{\theta}_t}(\mathbf{x}_u, \mathbf{A})$.

However, from the perspective of the computational overhead and practical performance, we only conduct dropout on the last layer during MC-dropout. In other words, the probability vector $\tilde{\mathbf{p}}_u^t \in \tilde{\mathbf{P}}^t = f_{\tilde{\theta}_t}(\mathbf{X}, \mathbf{A})$ at each time t can be obtained by:

$$\tilde{\mathbf{P}}^t = \sigma(\mathfrak{N}(\mathbf{A})(\mathbf{Z}^{(t)} \odot \sigma(\mathfrak{N}(\mathbf{A}) \cdots \sigma(\mathfrak{N}(\mathbf{A})\mathbf{X}\mathbf{W}^{(1)}) \cdots)\mathbf{W}^{(l-1)}))\mathbf{W}^{(l)}), \quad (8)$$

4.2.2 Dropedge Variational Inference. The dropedge variational inference takes a similar way with dropout variational inference, but imposes the randomness on the network topology instead.

Specifically, the l -th layer output of f_θ with dropedge can be written as:

$$\mathbf{H}^{(l)} = \sigma(\mathfrak{N}(\mathbf{A} \odot \mathbf{Z}^{(l-1)})\mathbf{H}^{(l-1)}\mathbf{W}^{(l-1)}), \quad (9)$$

where each element of $\mathbf{Z}^{(l)} \in \{0, 1\}^{|\mathcal{V}| \times |\mathcal{V}|}$ is also a sample of Bernoulli random variable, representing whether or not the corresponding edge in \mathbf{A} is removed.

Similarly, we only conduct dropedge on the last layer and perform T -times Monte-Carlo sampling (referred to as Monte-Carlo dropedge) base on dropedge, where at each time t , the probability vector $\tilde{\mathbf{p}}_u^t \in \tilde{\mathbf{P}}^t = f_{\tilde{\theta}_t}(\mathbf{X}, \mathbf{A})$ at each time t is obtained by

$$\tilde{\mathbf{P}}^t = \sigma(\mathfrak{N}(\mathbf{A} \odot \mathbf{Z}^{(t)})\sigma(\mathfrak{N}(\mathbf{A}) \cdots \sigma(\mathfrak{N}(\mathbf{A})\mathbf{X}\mathbf{W}^{(1)}) \cdots)\mathbf{W}^{(l-1)})\mathbf{W}^{(l)}), \quad (10)$$

4.2.3 Information Gain Estimation. With such probability vector $\tilde{\mathbf{p}}_u^t$ obtained by Eq. 8 or Eq. 10, we can calculate the prediction distribution $\mathbf{p}_u^{\mathcal{G}}$ by averaging all the $\tilde{\mathbf{p}}_u^t$:

$$\mathbf{p}_u^{\mathcal{G}} = \mathbf{p}(y_u | \mathbf{x}_u, \mathbf{A}, \mathcal{G}) = \frac{1}{T} \sum_{t=1}^T \tilde{\mathbf{p}}_u^t, \tilde{\theta}_t \sim P(\theta|\mathcal{G}), \quad (11)$$

and thus the information gain \mathbb{B}_u can be calculated by:

$$\mathbb{B}_u(y_u, \theta | \mathbf{x}_u, \mathbf{A}, \mathcal{G}) = - \sum_{d=1}^D p_{u,d}^{\mathcal{G}} \log p_{u,d}^{\mathcal{G}} + \frac{1}{T} \sum_{d=1}^D \sum_{t=1}^T \tilde{p}_{u,d}^t \log \tilde{p}_{u,d}^t. \quad (12)$$

Finally, we weight the loss function with above information gain after normalization:

$$\begin{aligned} \mathcal{L}_{st} &= \frac{|\mathcal{S}_U|}{|\mathcal{V}_L \cup \mathcal{S}_U|} \mathbb{E}_{(v_u, y_u) \sim P_{st}(\mathcal{V}, \mathcal{Y})} \bar{\mathbb{B}}_u l(\tilde{y}_u, \mathbf{p}_u) \\ &+ \frac{|\mathcal{V}_L|}{|\mathcal{V}_L \cup \mathcal{S}_U|} \mathbb{E}_{(v_i, y_i) \sim P_{pop}(\mathcal{V}, \mathcal{Y})} l(y_i, \mathbf{p}_i) \quad (13) \\ \text{where } \bar{\mathbb{B}}_u &= \frac{\mathbb{B}_u}{\beta \cdot \frac{1}{|\mathcal{S}_U|} \sum_i \mathbb{B}_i}. \end{aligned}$$

Here, we can tune the balance coefficient β to recover the population distribution (i.e., $\mathcal{L}_{st} \approx \mathcal{L}_{pop}$) as much as possible.

4.3 Improving Qualities of Pseudo Labels via Loss Correction

Till now, we have addressed the distribution shift issue with an information gain weighted loss function, where more attentions are paid to nodes with high information gain rather than high confidence. Unfortunately, such a training pipeline still implies hidden risks. Specifically, considering that pseudo labels of hard nodes are more likely to be incorrect as shown in Fig. 1(a) and our DR-GST focuses more on hard nodes, the impact of incorrect pseudo-labeled nodes will be enlarged and even mislead the learning of GCNs. Previous works generally filter out these low-quality nodes with collaborative scoring [20, 28] or prefabricated assumption [38] in a relatively coarse-grained manner, where abundant nodes with high information gain are discarded in advance. Instead, motivated by studies on learning with noisy labels [10, 23, 27], we propose to incorporate loss correction strategy into graph self-training. In brief, DR-GST corrects the predictions of the student model in each iteration, so as to eliminate the negative impact of misleading pseudo labels from the teacher model.

Specifically, given a student model $f_{\bar{\theta}}$ trained by pseudo labels, the loss correction assumes there is a model f_{θ^*} trained by ground-truth labels and a transition matrix \mathbf{T} such that $f_{\bar{\theta}}$ can be represented by $f_{\bar{\theta}} = \mathbf{T}f_{\theta^*}$, as shown in Fig. 3, where each element in $\mathbf{T} \in \mathbb{R}^{c \times c}$ is a transition probability from the ground-truth label to the pseudo label, i.e., $T_{kj} = P(\bar{Y} = j | Y = k)$ and c is the number of classes. With such a transition matrix, every model trained by pseudo labels is equal to that trained by ground-truth labels. We have proved the equivalence relation above using the following proposition.

PROPOSITION 4.2. *Given a model $f_{\bar{\theta}}$ trained by pseudo labels and a model f_{θ^*} trained by ground-truth labels, assuming that there exists a transition matrix \mathbf{T} such that the equation $f_{\bar{\theta}}(\mathbf{x}_u, \mathbf{A}) = \mathbf{T}f_{\theta^*}(\mathbf{x}_u, \mathbf{A})$ holds for each node v_u , then $\bar{\theta} = \theta^*$ if \mathbf{T} is a permutation matrix under cross entropy (CE) loss or \mathbf{T} is an arbitrary non-zero matrix under mean square error (MSE) loss.*

PROOF. Please refer to Appendix A.1.2. \square

Based on Proposition 4.2, ideally, we can train the student model regardless of the quality of labels, and recover f_{θ^*} with \mathbf{T} . Specifically, as shown in Fig. 3, for each node $v_i \in \{\mathcal{V}_L \cup \mathcal{S}_U\}$ with its feature vector \mathbf{x}_i , we first feed it into student model and multiply the output with \mathbf{T} to get $f_{\bar{\theta}}(\mathbf{x}_i, \mathbf{A})$. Then we use $f_{\bar{\theta}}(\mathbf{x}_i, \mathbf{A})$ to optimize the student model according to Eq. 13. Finally, at inference, we can treat the student model as f_{θ^*} . Please note that the transition matrix \mathbf{T} is pre-computed and not updated during optimization of the student model.




Figure 3: An illustration of loss correction.

Next, we make an illustration for the computation of the transition matrix \mathbf{T} . Noting that for each node $v_i \in \mathcal{V}_L$ with the ground-truth label $y_i = k$, the probability $P(Y = k | X = \mathbf{x}_i)$ should be 1 since we definitely know its label to be k . Therefore, given the output probability $p_{kj} = f_{\bar{\theta}}(\mathbf{x}_i, \mathbf{A})_j$ of class j , we have

$$\begin{aligned} p_{kj} &= P(\bar{Y} = j | X = \mathbf{x}_i) = \sum_{m=1}^c P(\bar{Y} = j | Y = m, X = \mathbf{x}_i) P(Y = m | X = \mathbf{x}_i) \\ &= P(\bar{Y} = j | Y = k, X = \mathbf{x}_i) \cdot 1 + 0 + \dots + 0 = T_{kj}(\mathbf{x}_i) = T_{kj}. \end{aligned} \quad (14)$$

In others words, the output probability vector $f_{\bar{\theta}}(\mathbf{x}_i, \mathbf{A})$ of each node v_i with its ground-truth label k is the k -th row of \mathbf{T} , where $\bar{\theta}$ means such a model is trained with the augmented dataset $\mathcal{V}_L \cup \mathcal{S}_U$.

Technically, we first train a student model $f_{\bar{\theta}}$ without loss correction using the augmented dataset $\mathcal{V}_L \cup \mathcal{S}_U$, then update \mathbf{T} according to $p_{kj} = f_{\bar{\theta}}(\mathbf{x}_i, \mathbf{A})_j$, and finally re-train a student model from scratch with loss correction to obtain f_{θ^*} .

Considering that there are multiple nodes belonging to class k in \mathcal{V}_L , we propose the following optimization problem to learn \mathbf{T} instead:

$$\arg \min_{\mathbf{T}} \sum_{k=1}^c \sum_{j=1}^{N_k^{(L)}} \|T_{k,:} - f_{\bar{\theta}}(\mathbf{x}_i, \mathbf{A})\|^2 + \|\mathbf{T}\mathbf{T}^T - \mathbf{I}\|^2, \quad (15)$$

where $N_k^{(L)}$ is the number of nodes belonging to class k in \mathcal{V}_L and \mathbf{I} is an identity matrix. Since the improved CE loss is utilized as the loss function in this paper as mentioned in Eq. 1 and Eq. 13, we append the regularization term $\|\mathbf{T}\mathbf{T}^T - \mathbf{I}\|^2$ for guiding \mathbf{T} to approximate to a permutation matrix, which is derived from Proposition 4.2 under the CE loss. Moreover, we initialize \mathbf{T} with the identify matrix \mathbf{I} at the very beginning.

4.4 Overview of DR-GST

Till now, we have elaborated our proposed DR-GST framework, which solves both the distribution shift and the low-quality pseudo labels with the help of information gain and loss correction. We summarize it in Algorithm 1 and further analyze its time complexity in Appendix A.2.

Given a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{X})$ with its original labeled dataset \mathcal{V}_L , unlabeled dataset \mathcal{V}_U , adjacent matrix \mathbf{A} as well as its label set \mathcal{Y}_L , we first train a teacher model $f_{\bar{\theta}}$ on \mathcal{V}_L to obtain the prediction \tilde{y}_u and the confidence r_u for each unlabeled node $v_u \in \mathcal{V}_U$ at line 1. then iterate steps from line 3 to 9 until convergence, where we call each iteration a *stage* following [28]. Specifically, at line 3 we select part of unlabeled nodes whose confidence r_u is bigger than a given threshold τ to obtain \mathcal{S}_U . Next at line 4 we pseudo-label each node $v_u \in \mathcal{S}_U$ with \tilde{y}_u to augment \mathcal{V}_L . Then at line 5 we calculate the information gain \mathbb{B}_u according to dropout or droppedge

variational inference in Section 4.2 and normalize it according to Eq. 13. With such information gain, we train a student model $f_{\hat{\theta}}$ at line 6 using the augmented dataset, where pseudo labels may be incorrect. Therefore, at line 7 we update the transition matrix T with the output probability vector of $f_{\hat{\theta}}$ of each node $v_i \in \mathcal{V}_L$ according to Eq. 15, and retrain the student model from scratch at line 8 with $f_{\hat{\theta}} = Tf_{\theta^*}$ to get f_{θ^*} . Finally, we replace the teacher model $f_{\hat{\theta}}$ with f_{θ^*} and repeat above steps until convergence.

Algorithm 1 The DR-GST Framework

Input: Graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{X})$, original labeled dataset \mathcal{V}_L , unlabeled dataset \mathcal{V}_U , adjacent matrix \mathbf{A} , label set \mathcal{Y}_L , transition matrix $T = \mathbf{I}$

Output: Probability vector \mathbf{p}_i for each node v_i

- 1: Train a teacher model $f_{\hat{\theta}}$ on \mathcal{V}_L to obtain the prediction \hat{y}_u and the confidence r_u for each unlabeled node $v_u \in \mathcal{V}_U$;
- 2: **for** each stage k **do**
- 3: Select part of unlabeled nodes according to r_u to get \mathcal{S}_U ;
- 4: Pseudo-labeling each node $v_u \in \mathcal{S}_U$ with \hat{y}_u ;
- 5: Calculate the information gain \mathbb{B}_u according to Eq. 12;
- 6: Train a student model $f_{\hat{\theta}}$ without T according to Eq. 13;
- 7: Update T using $f_{\hat{\theta}}(\mathbf{x}_i, \mathbf{A})$ of $v_i \in \mathcal{V}_L$ according to Eq. 15;
- 8: Retrain a student model from scratch according to Eq. 13 with $f_{\hat{\theta}} = Tf_{\theta^*}$ to get f_{θ^*} ;
- 9: Replace the teacher model $f_{\hat{\theta}}$ with the student model f_{θ^*} ;
- 10: **end for**
- 11: **return** $\mathbf{p}_i = f_{\theta^*}(\mathbf{x}_i, \mathbf{A})$ in the final stage.

4.5 Theoretical Analysis

In this section, we theoretically analyze the influence factors on self-training from the perspective of gradient descent, and our theorem below demonstrates the rationality of the whole DR-GST framework.

THEOREM 4.3. *Assuming that $\|\nabla_{\theta} l(y_i, \mathbf{p}_i)\| \leq \Psi$ for each node v_i , where Ψ is a constant, given $\nabla_{\theta} \mathcal{L}_{pop}$ and $\nabla_{\theta} \mathcal{L}_{st}$, the gradient of \mathcal{L}_{pop} and \mathcal{L}_{st} w.r.t. model parameters θ , the following bound between $\nabla_{\theta} \mathcal{L}_{pop}$ and $\nabla_{\theta} \mathcal{L}_{st}$ holds:*

$$\|\nabla_{\theta} \mathcal{L}_{pop} - \nabla_{\theta} \mathcal{L}_{st}\| \leq \frac{|\mathcal{S}_U|}{|\mathcal{V}_L \cup \mathcal{S}_U|} \Psi (2\|P_{(v_u, y_u) \sim P_{pop}(\mathcal{V}, \mathcal{Y})}(\bar{y}_u \neq y_u)\| + \|P_{st}(\mathcal{V}, \mathcal{Y}) - P_{pop}(\mathcal{V}, \mathcal{Y})\|). \quad (16)$$

PROOF. Please refer to Appendix A.1.3. \square

From the Theorem 4.3 we can conclude that the performance of self-training is negatively related to the difference $\|P_{st}(\mathcal{V}, \mathcal{Y}) - P_{pop}(\mathcal{V}, \mathcal{Y})\|$ between the two distributions as well as the error rate $\|P_{(v_u, y_u) \sim P_{pop}(\mathcal{V}, \mathcal{Y})}(\bar{y}_u \neq y_u)\|$ of pseudo labels. Meanwhile, we find our proposed DR-GST is a natural framework equipped with two designs to correspondingly address the issues in self-training: *information gain weighted loss function* for distribution recovery and *loss correction strategy* for improving qualities of pseudo labels. This analysis further demonstrates the rationality of DR-GST framework from the theoretical perspective.

5 EXPERIMENT

In this section, we evaluate the effectiveness of DR-GST framework on semi-supervised node classification task with five widely used benchmark datasets from citation networks [3, 26] (*i.e.*, Cora, Citeseer, Pubmed and CoraFull) and social networks [21] (*i.e.*, Flickr). More detailed descriptions about datasets are in Appendix A.3.1.

5.1 Experimental Setup

5.1.1 Baselines. We compare our proposed DR-GST framework with two categories of baselines, including three representative GCNs (*i.e.*, GCN [17], GAT [30], PPNP [18]) and three graph self-training frameworks (*i.e.*, STs [20], M3S [28], ABN [38]). Noting that STs includes four variants (*i.e.*, Self-Training, Co-Training, Union and Intersection) in the original paper and the best performance is reported in our experiments. The implementation of DR-GST and all the baselines can be seen in Appendix A.3.2. More detailed experimental environment can be seen in Appendix A.3.3.

5.1.2 Evaluation Protocol. To more comprehensively evaluate our model, for all the datasets, we arrange only a few (including 3, 5, 10, 20) labeled nodes per class (L/C) for the training set following [20]. Specifically, in the setting $L/C = 20$, we follow the standard split [26] for Cora, Citeseer and Pubmed, and manually select 20 labeled nodes per class for CoraFull and Flickr considering the lack of standard split. In the setting $L/C < 20$, we make 10 random splits for each L/C , where each random split represents that we randomly select part of nodes from the training set of $L/C = 20$. For all the methods and all the cases, we run 10 times and report the mean accuracy.

5.2 Overall Comparison on Node Classification

The performance of different methods on node classification are summarized in Table 1. We have the following observations.

- Our proposed DR-GST framework outperforms all the baselines by a considerable margin across most cases of all the datasets. The results demonstrate the effectiveness of DR-GST by adopting a more principled mechanism to make use of unlabeled nodes in graph for boosting classification performance.
- With the decrease of labeled nodes, we observe that the performance of GCNs (*i.e.*, GCN, GAT and APPNP) drops quickly. For clarity, we further illustrate the changing trend of accuracy *w.r.t.* L/C in Fig. 4. Obviously, we can discover the larger performance margin between DR-GST and GCNs with fewer labeled nodes per class, which further implies the superior capacity of DR-GST for addressing labeled data scarcity on graph learning.
- Considering the two variants of DR-GST, we find that DR-GST_{do} performs better on Pubmed, CoraFull and Flickr while DR-GST_{de} on Cora and Citeseer. An intuitive explanation for such distinct performance is the different emphasis on network topology and feature information *w.r.t.* different graphs for node classification task. Correspondingly, in DR-GST framework, MC-dropege performs information gain estimation with network topology while MC-dropout is based on feature information. This finding also sheds light on possible future work to combine both topology and feature to further enhance performance under our framework.

Table 1: Node classification results(%). (L/C: the number of labels per class; bold: best)

Dataset	Cora				Citeseer				Pubmed				CoraFull				Flickr			
L/C	3	5	10	20	3	5	10	20	3	5	10	20	3	5	10	20	3	5	10	20
GCN	64.52	69.55	78.03	81.56	51.39	61.34	68.39	71.64	66.04	71.25	75.88	79.31	41.83	49.12	55.67	60.69	37.69	40.64	48.04	51.74
GAT	67.19	69.45	76.38	82.24	55.19	59.40	67.61	72.00	67.85	68.41	72.42	78.38	36.44	46.70	52.45	57.97	20.02	24.90	33.27	37.06
APNP	65.06	75.53	81.33	83.14	51.22	60.48	68.50	71.64	65.77	73.01	76.35	79.51	40.29	44.49	50.89	60.77	24.76	35.54	47.87	61.55
STs	70.68	75.60	80.35	82.89	56.29	65.59	74.17	74.36	69.82	73.77	77.68	81.02	43.44	51.16	58.40	61.70	35.21	43.25	48.23	52.99
M3S	64.24	71.02	78.93	82.78	50.07	63.28	74.54	74.72	68.76	69.21	70.72	81.34	42.77	49.75	57.43	61.40	35.33	39.02	47.62	51.87
ABN	66.39	73.07	78.73	81.79	54.30	64.27	69.90	72.81	59.17	71.40	75.26	79.09	43.38	48.39	55.88	60.62	35.13	41.62	47.01	52.10
DR-GST _{do}	70.85	77.92	80.88	83.34	59.39	69.08	75.00	75.78	70.74	74.63	78.44	81.08	45.44	53.29	60.01	62.75	37.84	43.47	49.48	53.66
DR-GST _{de}	73.43	77.59	81.67	84.03	60.60	69.91	74.65	75.26	70.55	73.71	77.42	80.65	45.42	52.50	59.16	63.11	38.21	43.28	49.44	53.05

- Among the two categories of baselines, self-training frameworks (*i.e.*, STs, M3S and ABN) can generally improve GCNs (*i.e.*, GCN, GAT and APPNP), which indicates the usefulness of unlabeled data. Nevertheless, DR-GST still yields better performance for the following two promising designs: 1) We pay more attention on nodes with high information gain rather than high confidence, so that the unsatisfying distribution shift issue is avoided. 2) We adopt a loss correction strategy, where qualities of pseudo labels are improved for subsequent self-training.

5.3 In-depth Analysis of DR-GST

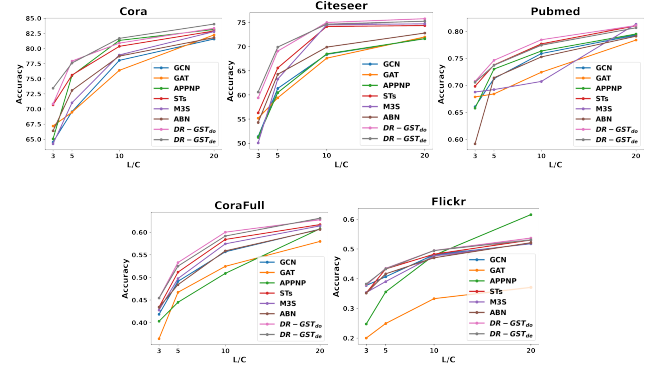
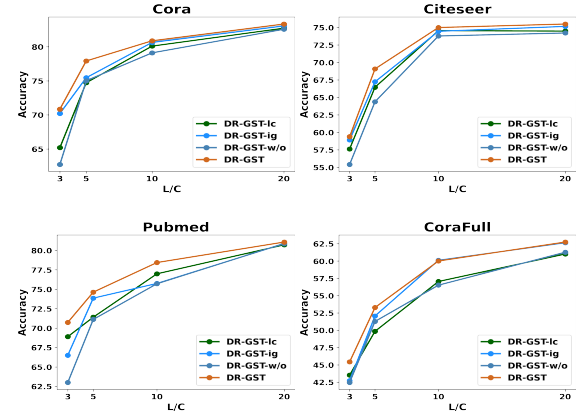
In this section, we make a series of analysis to better understand each component in DR-GST, as well as key parameter selections.

5.3.1 Ablation Study. As mentioned above, the performance of self-training theoretically hinges on the distribution gap and qualities of pseudo labels, which could be naturally captured by our DR-GST framework with two corresponding designs: the information-gain based weighted loss function and loss correction module. To comprehensively understand their contributions towards self-training on graphs, we prepare following three variants of DR-GST:

- **DR-GST-lc:** DR-GST only with the loss correction module, *i.e.*, $\mathbb{B} = 1$ for all the unlabeled nodes.
- **DR-GST-ig:** DR-GST only with the information gain weighted loss function.
- **DR-GST-w/o:** DR-GST without the above two designs.

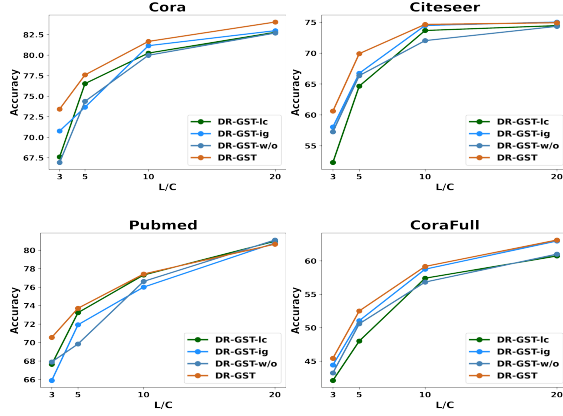


The results on DR-GST_{do} and DR-GST_{de} are respectively reported in Fig. 5 and Fig. 6. From the results we can find that the overall performance order is as follows: DR-GST > DR-GST-ig > DR-GST-lc > DR-GST-w/o. There are three conclusions here. Firstly, the best performance achieved by the complete DR-GST framework indicates the effectiveness of considering two components together. Secondly, the information gain weighted loss function and loss correction are both value modules for self-training on graphs. Thus, ignoring them altogether (*i.e.*, DR-GST-w/o) is not ideal. Thirdly, the information-gain weighted loss function plays a more vital role in our self-training framework since DR-GST-lc generally does not perform as well as DR-GST-ig. In short, above findings further verify the rationality of DR-GST from the empirical perspective.

5.3.2 Parameter Study. Here, we investigate into the sensitivity of two hyper-parameters (*i.e.*, threshold τ and balance coefficient β) on Cora and CoraFull datasets. Similar observations are also

**Figure 4: The changing trends of accuracy w.r.t. L/C****Figure 5: Ablation study of DR-GST_{do}.**

made on other datasets. In particular, we respectively report the performance of DR-GST_{do} and DR-GST_{de}, and vary the L/C in {3, 5, 10}. For clear notation in figures, we use “do-3” to denote DR-GST_{do} with L/C = 3, and the rest can be done in the same manner.

Analysis of threshold τ in self-training We test the impact of threshold τ in self-training, and vary it from 0.40 to 0.70 for Cora and 0.60 to 0.90 for CoraFull. The results are summarized in Fig

Figure 6: Ablation study of DR-GST_{de}.Figure 7: Impact of threshold τ .Figure 8: Impact of balance coefficient β .

.7. Generally speaking, the best performance is achieved when we set a smaller τ , which is consistent with our analysis above that high-confidence unlabeled nodes contribute less.

Analysis of balance coefficient β We then test the impact of the balance coefficient β in Eq. 13, and vary it from 1/3 to 1. The results are shown in Fig. 8. Obviously, with the increase of β , or, in other words, with more attention paid to hard nodes, the performance shows a downward trend, further demonstrating the effectiveness of our design.

5.3.3 Visualization. For a more intuitive of the proposed information gain based DR-GST, we conduct the task of visualization on Cora dataset. Specifically, as shown in Fig. 9, we visualize the output embedding of the student model at different stages in DR-GST for Cora dataset. From Fig. 9(a) to Fig. 9(c) we show the visualization of unlabeled nodes, where a lighter dot represents a node endowed with a higher weight by information gain when calculating the loss

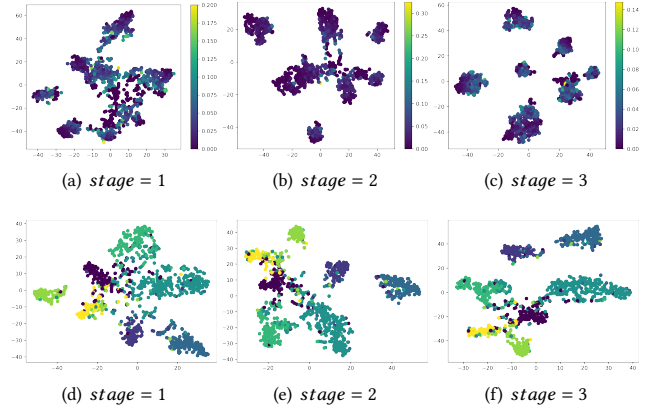


Figure 9: Visualization of learned embeddings for unlabeled nodes ((a)~(c)) and test nodes ((d)~(f)) on Cora at different stages during self-training.

function in Eq. 13. Obviously, we can discover that at an earlier stage, DR-GST pays more attention to nodes close to the decision boundary which is also indistinct at this moment. With the training progress going on, the light nodes gradually vanish, implying that most of information these nodes contain has been learned, leading to a more crisp decision boundary. From Fig. 9(d) to Fig. 9(f) we show the visualization of test nodes, where different colors represent different classes. Apparently, the separability of different classes for test nodes is gradually improved, further demonstrating the effectiveness of DR-GST for optimizing the decision boundary.

6 RELATED WORK

In line with the main focus of our work, we review the most related work in graph neural networks and self-training.

Graph Neural Networks Recent years have seen a surge of efforts on Graph Neural Networks (GNNs) and achieved state-of-the-art performance in various tasks on graphs [33, 36]. Generally, current GNNs can be divided into two categories. The first category is spectral-based GNNs, which defines graph convolution operation in the spectral domain [4, 7]. The well-known GCN [17] simplifies graph convolutions by using the 1-order approximation. Since then, plenty of studies have sprung up. SGC [32] further simplifies GCN by removing the nonlinearities between GCN layers. [20] shows that GCNs smooth node features between neighbours. On the comparison, the other category is spatial-based GNNs, mainly devoted to aggregating and transforming the local information from the perspective of spatial domain. GAT [30] assigns the learnt weight to each edge during aggregation. [12] proposes a permutation-invariant aggregator for message passing. Moreover, there are many other graph neural models, we please refer the readers to recent surveys [34, 37] for a more comprehensive review.

Self-training Despite the success, GNNs typically require large amounts of labeled data, which is expensive and time-consuming. Self-training [16] is one of the earliest strategies addressing labeled data scarcity by making better use of abundant unlabeled data, and has shown remarkable performance on various tasks [14, 19, 22].

Recently, [20] proposes a graph-based self-training framework, demonstrating the effectiveness of self-training on graphs. Further, [28] utilizes the DeepCluster [5] to filter out low-quality pseudo labels during self-training. CaGCN-st [31] argues that self-training under-performs due to generally overlooked low-confidence but high-accuracy predictions, and proposes a confidence-calibrated self-training framework. [38] proposes to select high-quality unlabeled nodes via an adaptive pseudo labeling technique. [24] utilizes a margin prediction confidence to select unlabeled nodes, aiming at identifying the most confident labels. In summary, almost all of graph self-training methods focus on improving the quality of pseudo labels by virtue of confidence, but none of them have ever considered the capability and limitation of such selection criterion.

7 CONCLUSION

In this paper, we empirically make a thorough study for capability and limitation of current self-training methods on graphs, and surprisingly find they may be cheated by confidence and even suffer from the distribution shift issue, leading to unpromising performance. To this end, we propose a novel self-training framework DR-GST which not only addresses the distribution shift issue from the view of information gain, but also is equipped with the creative loss correction strategy for improving qualities of pseudo labels. Theoretical analysis and extensive experiments well demonstrate the effectiveness of the proposed DR-GST. Moreover, our study also gives an insight that confidence alone is not enough for self-training and thus motivates us an interesting direction for future work, *i.e.*, exploiting more criteria for the selection of unlabeled nodes during self-training.

8 ACKNOWLEDGMENTS

This work is supported in part by the National Natural Science Foundation of China (No. U20B2045, 62192784, 62172052, 61772082, 62002029, U1936104), the Fundamental Research Funds for the Central Universities 2021RC28 and CCF-Ant Group Research Fund.

REFERENCES

- [1] Antonin Berthon, Bo Han, Gang Niu, Tongliang Liu, and Masashi Sugiyama. 2021. Confidence scores make instance-dependent label-noise learning possible. In *ICML*. 825–836.
- [2] Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. 2015. Weight uncertainty in neural network. In *ICML*. 1613–1622.
- [3] Aleksandar Bojchevski and Stephan Günnemann. 2018. Deep Gaussian Embedding of Graphs: Unsupervised Inductive Learning via Ranking. In *ICLR*.
- [4] Joan Bruna, Wojciech Zaremba, Arthur Szlam, and Yann LeCun. 2014. Spectral Networks and Locally Connected Networks on Graphs. In *ICLR*.
- [5] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. 2018. Deep clustering for unsupervised learning of visual features. In *ECCV*. 132–149.
- [6] Thomas M Cover. 1999. *Elements of information theory*. John Wiley & Sons.
- [7] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. 2016. Convolutional Neural Networks on Graphs with Fast Localized Spectral Filtering. In *NIPS*. 3837–3845.
- [8] Jiali Duan, Yen-Liang Lin, Son Dinh Tran, Larry S. Davis, and C.-C. Jay Kuo. 2020. SLADE: A Self-Training Framework for Distance Metric Learning. In *CVPR*. 9644–9653.
- [9] Yarín Gal and Zoubin Ghahramani. 2016. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *ICML*. 1050–1059.
- [10] Jacob Goldberger and Ehud Ben-Reuven. 2017. Training deep neural-networks using a noise adaptation layer. In *ICLR*.
- [11] Alex Graves. 2011. Practical variational inference for neural networks. In *NIPS*. 2348–2356.
- [12] William L Hamilton, Rex Ying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. In *NIPS*. 1025–1035.
- [13] Arman Hasanzadeh, Ehsan Hajiramezani, Shahin Boluki, Mingyuan Zhou, Nick Duffield, Krishna Narayanan, and Xiaoning Qian. 2020. Bayesian graph neural networks with adaptive connection sampling. In *ICML*. 4094–4104.
- [14] Junxian He, Jiatao Gu, Jiajun Shen, and Marc Aurelio Ranzato. 2020. Revisiting Self-Training for Neural Sequence Generation. In *ICLR*.
- [15] Jose Hernandez-Lobato, Yingzhen Li, Mark Rowland, Thang Bui, Daniel Hernández-Lobato, and Richard Turner. 2016. Black-box alpha divergence minimization. In *ICML*. 1511–1520.
- [16] H. J. Scudder III. 1965. Probability of error of some adaptive pattern-recognition machines. *IEEE Transactions on Information Theory* 11, 3 (1965), 363–371.
- [17] Thomas N. Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *ICLR*.
- [18] Johannes Klicpera, Aleksandar Bojchevski, and Stephan Günnemann. 2019. Predict then Propagate: Graph Neural Networks meet Personalized PageRank. In *ICLR*.
- [19] Dong-Hyun Lee et al. 2013. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *ICML Workshop*.
- [20] Qimai Li, Zhichao Han, and Xiao-Ming Wu. 2018. Deeper Insights Into Graph Convolutional Networks for Semi-Supervised Learning. In *AAAI*. 3538–3545.
- [21] Zaïqiao Meng, Shangsong Liang, Hongyan Bao, and Xiangliang Zhang. 2019. Co-embedding attributed networks. In *WSDM*. 393–401.
- [22] Subhabrata Mukherjee and Ahmed Awadallah. 2020. Uncertainty-aware self-training for few-shot text classification. In *NIPS*.
- [23] Giorgio Patrini, Alessandro Rozza, Aditya Krishna Menon, Richard Nock, and Lizhen Qu. 2017. Making deep neural networks robust to label noise: A loss correction approach. In *CVPR*. 1944–1952.
- [24] Daniel Carlos Guimarães Pedronette and Longin Jan Latecki. 2021. Rank-based self-training for graph convolutional networks. *Information Processing & Management* 58, 2 (2021), 102443.
- [25] Yu Rong, Wenbing Huang, Tingyang Xu, and Junzhou Huang. 2019. Droppedge: Towards deep graph convolutional networks on node classification. *arXiv preprint arXiv:1907.10903* (2019).
- [26] Prithviraj Sen, Galileo Namata, Mustafa Bilgic, Lise Getoor, Brian Galligher, and Tina Eliassi-Rad. 2008. Collective classification in network data. *AI magazine* 29, 3 (2008), 93–93.
- [27] Sainbayar Sukhbaatar, Joan Bruna, Manohar Paluri, Lubomir Bourdev, and Rob Fergus. 2014. Training convolutional networks with noisy labels. *arXiv preprint arXiv:1406.2080* (2014).
- [28] Ke Sun, Zhouchen Lin, and Zhanxing Zhu. 2020. Multi-Stage Self-Supervised Learning for Graph Convolutional Networks on Graphs with Few Labeled Nodes. In *AAAI*. 5892–5899.
- [29] Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of machine learning research* 9, 11 (2008).
- [30] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. 2017. Graph attention networks. *ICLR*.
- [31] Xiao Wang, Hongrui Liu, Chuan Shi, and Cheng Yang. 2021. Be Confident! Towards Trustworthy Graph Neural Networks via Confidence Calibration. *Advances in Neural Information Processing Systems* 34 (2021).
- [32] Felix Wu, Amauri Souza, Tianyi Zhang, Christopher Fifty, Tao Yu, and Kilian Weinberger. 2019. Simplifying graph convolutional networks. In *ICML*. 6861–6871.
- [33] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip. 2020. A comprehensive survey on graph neural networks. *IEEE transactions on neural networks and learning systems* 32, 1 (2020), 4–24.
- [34] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip. 2020. A comprehensive survey on graph neural networks. *IEEE transactions on neural networks and learning systems* 32, 1 (2020), 4–24.
- [35] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. 2019. How Powerful are Graph Neural Networks. In *ICLR*.
- [36] Ziwei Zhang, Peng Cui, and Wenwu Zhu. 2020. Deep learning on graphs: A survey. *IEEE Transactions on Knowledge and Data Engineering* (2020).
- [37] Jie Zhou, Ganqu Cui, Shengding Hu, Zhengyan Zhang, Cheng Yang, Zhiyuan Liu, Lifeng Wang, Changcheng Li, and Maosong Sun. 2020. Graph neural networks: A review of methods and applications. *AI Open* 1 (2020), 57–81.
- [38] Ziang Zhou, Jieming Shi, Shengzhong Zhang, Zengfeng Huang, and Qing Li. 2019. Effective Semi-Supervised Node Classification on Few-Labeled Graph Data. *arXiv preprint arXiv:1910.02684* (2019).

A SUPPLEMENT

In the supplement, we first provide detailed proofs of import theorems in our paper *i.e.*, Theorem 4.1, Proposition 4.2 and Theorem 4.3. Next, more experimental details are represented for reproduction.

A.1 Proof

In this section, we successively show the detailed proof for Theorem 4.1, Proposition 4.2 and Theorem 4.3.

A.1.1 Proof of Theorem 4.1.

PROOF. With our assumption that $\tilde{y}_u = y_u$ for each pseudo-labeled node $v_u \in \mathcal{S}_U$, we first rewrite \mathcal{L}_{pop} in Eq. 3 as:

$$\begin{aligned} \mathcal{L}_{pop} &= \frac{|\mathcal{S}_U|}{|\mathcal{V}_L \cup \mathcal{S}_U|} \mathbb{E}_{(v_u, y_u) \sim P_{pop}(\mathcal{V}, \mathcal{Y})} l(\tilde{y}_u, \mathbf{p}_u) \\ &\quad + \frac{|\mathcal{V}_L|}{|\mathcal{V}_L \cup \mathcal{S}_U|} \mathbb{E}_{(v_i, y_i) \sim P_{pop}(\mathcal{V}, \mathcal{Y})} l(y_i, \mathbf{p}_i). \end{aligned} \quad (17)$$

Note that

$$\mathbb{E}_{(v_u, y_u) \sim P_{pop}(\mathcal{V}, \mathcal{Y})} l(\tilde{y}_u, \mathbf{p}_u) = \mathbb{E}_{(v_u, y_u) \sim P_{st}(\mathcal{V}, \mathcal{Y})} \frac{P_{pop}(v_u, y_u)}{P_{st}(v_u, y_u)} l(\tilde{y}_u, \mathbf{p}_u), \quad (18)$$

then we can rewrite Eq. 17 as

$$\begin{aligned} \mathcal{L}_{pop} &= \frac{|\mathcal{S}_U|}{|\mathcal{V}_L \cup \mathcal{S}_U|} \mathbb{E}_{(v_u, y_u) \sim P_{st}(\mathcal{V}, \mathcal{Y})} \frac{P_{pop}(v_u, y_u)}{P_{st}(v_u, y_u)} l(\tilde{y}_u, \mathbf{p}_u) \\ &\quad + \frac{|\mathcal{V}_L|}{|\mathcal{V}_L \cup \mathcal{S}_U|} \mathbb{E}_{(v_i, y_i) \sim P_{pop}(\mathcal{V}, \mathcal{Y})} l(y_i, \mathbf{p}_i) \\ &= \frac{|\mathcal{S}_U|}{|\mathcal{V}_L \cup \mathcal{S}_U|} \mathbb{E}_{(v_u, y_u) \sim P_{st}(\mathcal{V}, \mathcal{Y})} \gamma_u l(\tilde{y}_u, \mathbf{p}_u) \\ &\quad + \frac{|\mathcal{V}_L|}{|\mathcal{V}_L \cup \mathcal{S}_U|} \mathbb{E}_{(v_i, y_i) \sim P_{pop}(\mathcal{V}, \mathcal{Y})} l(y_i, \mathbf{p}_i), \end{aligned} \quad (19)$$

where γ_u can be regarded as a weight of the loss function for each pseudo-labeled node v_u .

Finally, recalling the loss function under the distribution shift case in Eq. 4, *i.e.*,

$$\begin{aligned} \mathcal{L}_{st} &= \frac{|\mathcal{V}_L|}{|\mathcal{V}_L \cup \mathcal{S}_U|} \mathbb{E}_{(v_i, y_i) \sim P_{pop}(\mathcal{V}, \mathcal{Y})} l(y_i, \mathbf{p}_i) \\ &\quad + \frac{|\mathcal{S}_U|}{|\mathcal{V}_L \cup \mathcal{S}_U|} \mathbb{E}_{(v_u, y_u) \sim P_{st}(\mathcal{V}, \mathcal{Y})} l(\tilde{y}_u, \mathbf{p}_u), \end{aligned} \quad (20)$$

we can find that it is definitely equal to that in Eq. 3 with an additional weight coefficient. In other words, we can recover the population distribution as long as we weight each pseudo-labeled node with a proper coefficient in \mathcal{L}_{st} . \square

A.1.2 Proof of Proposition 4.2.

PROOF. Without loss of generality, we respectively prove the equality of θ^* and $\bar{\theta}$ under MSE loss and CE loss.

MSE loss. Under the MSE loss, with our non-zero assumption for \mathbf{T} , the following equation holds true:

$$\begin{aligned} \theta^* &= \arg \min_{\theta^* \in \Theta} \sum_u \|f_{\theta^*}(\mathbf{x}_u, \mathbf{A}) - \mathbf{y}_u\|^2 \\ &= \arg \min_{\theta^* \in \Theta} \sum_u \|\mathbf{T}f_{\theta^*}(\mathbf{x}_u, \mathbf{A}) - \mathbf{T}\mathbf{y}_u\|^2 \\ &= \arg \min_{\bar{\theta} \in \Theta} \sum_u \|f_{\bar{\theta}}(\mathbf{x}_u, \mathbf{A}) - \tilde{\mathbf{y}}_u\|^2 = \bar{\theta}, \end{aligned} \quad (21)$$

where \mathbf{y}_u is a one-hot vector expanded from y_u . The proof is concluded for MSE loss.

CE loss. Under CE loss, we prove the equality of $\bar{\theta}$ and θ^* from the perspective of gradient descent. Specifically, if for each node v_u , the gradient of $f_{\bar{\theta}}(\mathbf{x}_u, \mathbf{A})$ w.r.t. $\bar{\theta}$ is equal to that of $f_{\theta^*}(\mathbf{x}_u, \mathbf{A})$ w.r.t. θ^* , then optimizing a model $f_{\bar{\theta}}$ using gradient descent will definitely leads to our desired model f_{θ^*} , that is to say, $\bar{\theta} = \theta^*$.

Specifically, for each node v_u , we first rewrite the CE loss as follows:

$$l(\mathbf{y}_u, \mathbf{p}_u) = \mathbf{y}_u^T \log f_{\bar{\theta}}(\mathbf{x}_u, \mathbf{A}). \quad (22)$$

Then the difference d of gradient between $\bar{\theta}$ and θ^* can be written as:

$$d = \|\nabla_{\bar{\theta}} \mathbf{y}_u^T \log f_{\bar{\theta}}(\mathbf{x}_u, \mathbf{A}) - \nabla_{\theta^*} \mathbf{y}_u^T \log f_{\theta^*}(\mathbf{x}_u, \mathbf{A})\| \quad (23)$$

Considering our assumption that $f_{\bar{\theta}}(\mathbf{x}_u, \mathbf{A}) = \mathbf{T}f_{\theta^*}(\mathbf{x}_u, \mathbf{A})$, Eq. 23 becomes:

$$d = \|\nabla_{\bar{\theta}} (\mathbf{T}\mathbf{y}_u)^T \log (\mathbf{T}f_{\theta^*}(\mathbf{x}_u, \mathbf{A})) - \nabla_{\theta^*} \mathbf{y}_u^T \log f_{\theta^*}(\mathbf{x}_u, \mathbf{A})\| \quad (24)$$

According to the chain rule, we have:

$$d = \|\nabla_{\theta^*} f_{\theta^*}(\mathbf{x}_u, \mathbf{A}) \cdot (\mathbf{T}^T (\mathbf{T}\mathbf{y}_u \oslash \mathbf{T}f_{\theta^*}(\mathbf{x}_u, \mathbf{A})) - \mathbf{y}_u \oslash f_{\theta^*}(\mathbf{x}_u, \mathbf{A}))\|, \quad (25)$$

where \oslash represents the element-wise division operation.

Obviously, if \mathbf{T} is a permutation matrix, the difference d of gradient is zero. The proof is concluded for CE loss. \square

A.1.3 Proof of Theorem 4.3. To prove Theorem 4.3, we need to borrow a corollary from [38], which illustrates the impact of incorrect pseudo labels on self-training without distribution shift.

COROLLARY A.1. Assuming that the augmented dataset follows the population distribution P_{pop} and $\|\nabla_{\theta} l\| \leq \Psi$ for any gradient $\nabla_{\theta} \mathcal{L}$, the following bound between $\nabla_{\theta} \mathcal{L}_{pop}$ and $\nabla_{\theta} \mathcal{L}_{st}$ holds:

$$\|\nabla_{\theta} \mathcal{L}_{pop} - \nabla_{\theta} \mathcal{L}_{st}\| \leq \frac{|\mathcal{S}_U|}{|\mathcal{V}_L \cup \mathcal{S}_U|} 2\Psi \|P_{(v_u, y_u) \sim P_{pop}(\mathcal{V}, \mathcal{Y})}(\tilde{y}_u \neq y_u)\|. \quad (26)$$

Now, we prove Theorem 4.3.

PROOF. We first calculate the difference between $\nabla_{\theta} \mathcal{L}_{pop}$ and $\nabla_{\theta} \mathcal{L}_{st}$ as follows:

$$\begin{aligned} \|\nabla_{\theta} \mathcal{L}_{pop} - \nabla_{\theta} \mathcal{L}_{st}\| &= \frac{|\mathcal{S}_U|}{|\mathcal{V}_L \cup \mathcal{S}_U|} \|\mathbb{E}_{(v_u, y_u) \sim P_{pop}(\mathcal{V}, \mathcal{Y})} \nabla_{\theta} l(y_u, \mathbf{p}_u) \\ &\quad - \mathbb{E}_{(v_u, y_u) \sim P_{st}(\mathcal{V}, \mathcal{Y})} \nabla_{\theta} l(\tilde{y}_u, \mathbf{p}_u)\|. \end{aligned} \quad (27)$$

Adding and subtracting a same term $\mathbb{E}_{(v_u, y_u) \sim P_{st}(\mathcal{V}, \mathcal{Y})} \nabla_{\theta} l(\tilde{y}_u, \mathbf{p}_u)$, and abbreviating $\frac{|\mathcal{S}_U|}{|\mathcal{V}_L \cup \mathcal{S}_U|}$ as η , Eq. 27 can be written as:

$$\begin{aligned} \|\nabla_{\theta} \mathcal{L}_{pop} - \nabla_{\theta} \mathcal{L}_{st}\| &= \\ \eta &\|\mathbb{E}_{(v_u, y_u) \sim P_{pop}(\mathcal{V}, \mathcal{Y})} \nabla_{\theta} l(y_u, \mathbf{p}_u) - \mathbb{E}_{(v_u, y_u) \sim P_{pop}(\mathcal{V}, \mathcal{Y})} \nabla_{\theta} l(\tilde{y}_u, \mathbf{p}_u) \\ &\quad + \mathbb{E}_{(v_u, y_u) \sim P_{pop}(\mathcal{V}, \mathcal{Y})} \nabla_{\theta} l(\tilde{y}_u, \mathbf{p}_u) - \mathbb{E}_{(v_u, y_u) \sim P_{st}(\mathcal{V}, \mathcal{Y})} \nabla_{\theta} l(\tilde{y}_u, \mathbf{p}_u)\|. \end{aligned} \quad (28)$$

According to the triangle property of the norm, the following inequality is satisfied:

$$\begin{aligned} \|\nabla_{\theta} \mathcal{L}_{pop} - \nabla_{\theta} \mathcal{L}_{st}\| &\leq \\ \eta &(\|\mathbb{E}_{(v_u, y_u) \sim P_{pop}(\mathcal{V}, \mathcal{Y})} \nabla_{\theta} l(y_u, \mathbf{p}_u) - \mathbb{E}_{(v_u, y_u) \sim P_{pop}(\mathcal{V}, \mathcal{Y})} \nabla_{\theta} l(\tilde{y}_u, \mathbf{p}_u)\| \\ &\quad + \|\mathbb{E}_{(v_u, y_u) \sim P_{pop}(\mathcal{V}, \mathcal{Y})} \nabla_{\theta} l(\tilde{y}_u, \mathbf{p}_u) - \mathbb{E}_{(v_u, y_u) \sim P_{st}(\mathcal{V}, \mathcal{Y})} \nabla_{\theta} l(\tilde{y}_u, \mathbf{p}_u)\|). \end{aligned} \quad (29)$$

Table 2: The statistics of datasets

Dataset	Nodes	Edges	Classes	Features	Validation	Test
Cora	2708	5429	7	1433	500	1000
Citeseer	3327	4732	6	3703	500	1000
Pubmed	19717	44338	3	500	500	1000
CoraFull	19793	65311	70	8710	500	1000
Flickr	7575	239738	9	12047	500	1000

Recalling Corollary A.1, we know that the first term on the right hand side satisfies:

$$\begin{aligned} & \|\mathbb{E}_{(v_u, y_u) \sim P_{pop}(\mathcal{V}, \mathcal{Y})} \nabla_{\theta} l(y_u, \mathbf{p}_u) - \mathbb{E}_{(v_u, y_u) \sim P_{pop}(\mathcal{V}, \mathcal{Y})} \nabla_{\theta} l(\tilde{y}_u, \mathbf{p}_u)\| \\ & \leq 2\Psi \|P_{(v_u, y_u) \sim P_{pop}(\mathcal{V}, \mathcal{Y})}(\tilde{y}_u \neq y_u)\|. \end{aligned} \quad (30)$$

And for the second term, we have:

$$\begin{aligned} & \|\mathbb{E}_{(v_u, y_u) \sim P_{pop}(\mathcal{V}, \mathcal{Y})} \nabla_{\theta} l(\tilde{y}_u, \mathbf{p}_u) - \mathbb{E}_{(v_u, y_u) \sim P_{st}(\mathcal{V}, \mathcal{Y})} \nabla_{\theta} l(\tilde{y}_u, \mathbf{p}_u)\| \\ & = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \nabla_{\theta} l(\tilde{y}_u, \mathbf{p}_u) d(P_{pop}(\mathcal{V}, \mathcal{Y}) - P_{st}(\mathcal{V}, \mathcal{Y})) \\ & \leq \Psi \cdot \|P_{pop}(\mathcal{V}, \mathcal{Y}) - P_{st}(\mathcal{V}, \mathcal{Y})\|, \end{aligned} \quad (31)$$

where the inequality is from our assumption that $\|\nabla_{\theta} l\| \leq \Psi$.

Combining Eq. 30 with Eq. 31, we have:

$$\begin{aligned} \|\nabla_{\theta} \mathcal{L}_{pop} - \nabla_{\theta} \mathcal{L}_{st}\| & \leq \frac{|S_U|}{|\mathcal{V}_L \cup S_U|} \Psi (2\|P_{(v_u, y_u) \sim P_{pop}(\mathcal{V}, \mathcal{Y})}(\tilde{y}_u \neq y_u) \\ & + \|P_{st}(\mathcal{V}, \mathcal{Y}) - P_{pop}(\mathcal{V}, \mathcal{Y})\|). \end{aligned} \quad (32)$$

The proof is concluded. \square

A.2 Time Complexity Analysis

We first analyze the time complexity of a general self-training framework. Assuming training an epoch takes $O(M)$ time, given epochs E , its time complexity in each stage is $O(EM)$. DR-GST is innovated in information gain and loss correction, which respectively takes $O(TM)$ and $O(Ec^2)$ time in each stage, where T and c are the numbers of sampling for variational inference and class. Moreover, considering that we train a student model twice in each stage, the total time complexity is $O((2E + T)M + Ec^2)$. In fact, T and $O(Ec^2)$ are always far less than E and $O(EM)$. Consequently, the time complexity of DR-GST is approximately twice that of the general self-training framework.

A.3 More Experimental Details

A.3.1 Details of datasets. We adopt five widely used benchmark datasets from citation networks [3, 26] (*i.e.*, Cora, Citeseer, Pubmed and CoraFull) and social network [21] (*i.e.*, Flickr) for evaluation. For the citation networks, nodes represent papers, edges are the citation relationship between papers, node features are comprised of bag-of-words vector of the papers and labels represent the fields of papers. And for the social network, nodes in Flickr represent users of the Flickr website, edges are their relationships induced by their photo-sharing records and labels represent users' interest groups. For all the datasets, We choose 500 nodes for validation, 1000 nodes for test. The details of these datasets are summarized in Table 2. Our data

are public and do not contain personally identifiable information and offensive content. The address of our data is <https://docs.dgl.ai/en/latest/api/python/dgl.data.html#node-prediction-datasets> and the license is Apache License 2.0.

A.3.2 Implementation. We supplement the implementation details of DR-GST and all the baselines here.

For fair comparison, we utilize the standard GCN with 2 layers as the backbone for all graph self-training framework. We optimize models via Adam with learning rate of 0.01 and early stopping with a window size of 200. In particular, we set L2 regularization with $\lambda_r = 5e - 4$ for Cora, Citeseer, Pubmed, CoraFull and $\lambda_r = 5e - 5$ for Flickr. We set ReLU as the activation function and apply a dropout rate of 0.5 to prevent over-fitting. As for the MC-dropout and MC-droppedge, we set the number of sampling $T = 100$. Moreover, we apply grid search for other important hyper-parameters. Specifically, the drop rate of MC-dropout and MC-droppedge is chosen from $\{0.1, 0.2, \dots, 0.5\}$, the balance coefficient β for information gain in Eq. 13 is searched in $\{4/3, 1, 2/3, 1/2, 1/3, 1/4\}$ and the threshold τ is tuned amongst $\{0.4, 0.45, \dots, 0.75\}$ for Cora, Citeseer, $\{0.6, 0.65, \dots, 0.9\}$ for Pubmed, CoraFull and $\{0.75, 0.78, \dots, 0.96\}$ for Flickr.

We adopt the implementation of GCN, GAT and APPNP from DGL², and the implementations of STs³ and ABN⁴ are publicly provided by their authors. Considering that the implementation of M3S is not available, we re-implement it referring to the original paper [28]. For all baselines, we perform grid search for important hyper-parameters (*i.e.*, τ) to obtain optimal results.

A.3.3 Experimental Environment. In this section we summarize the hardware and software environment in our experiments.

We utilize a linux machine powered by an Intel(R) Xeon(R) CPU E5-2682 v4 @ 2.50GHz CPU and 4 Tesla P100-PCIE-16GB as well as 4 GeForce RTX 3090 GPU cards.

The operating system is Linux version 3.10.0-693.el7.x86_64. We realize our code with Python 3.8.8 as well as some other python packages as follows: PyTorch 1.8.1, DGL 0.6.0 (cuda 10.1), NetworkX 2.5.

²<https://www.dgl.ai/>

³https://github.com/Davidham3/deeper_insights_into_GCNS

⁴<https://anonymous.4open.science/r/e7aca211-0d8d-4564-8f3f-0ef24b01941e/>