

Robust Training of Graph Neural Networks via Noise Governance

Siyi Qian
Zhejiang University
syqian99@zju.edu.cn

Haochao Ying*
Zhejiang University
haochaoying@zju.edu.cn

Renjun Hu
Alibaba Group
renjun.hrj@alibaba-inc.com

Jingbo Zhou
Baidu Research
zhoujingbo@outlook.com

Jintai Chen
Zhejiang University
jtchen721@gmail.com

Danny Z. Chen
University of Notre Dame
dchen@nd.edu

Jian Wu
Zhejiang University
wujian2000@zju.edu.cn

ABSTRACT

Graph Neural Networks (GNNs) have become widely-used models for semi-supervised learning. However, the robustness of GNNs in the presence of label noise remains a largely under-explored problem. In this paper, we consider an important yet challenging scenario where labels on nodes of graphs are not only noisy but also scarce. In this scenario, the performance of GNNs is prone to degrade due to label noise propagation and insufficient learning. To address these issues, we propose a novel RTGNN (Robust Training of Graph Neural Networks via Noise Governance) framework that achieves better robustness by learning to explicitly govern label noise. More specifically, we introduce self-reinforcement and consistency regularization as supplemental supervision. The self-reinforcement supervision is inspired by the memorization effects of deep neural networks and aims to correct noisy labels. Further, the consistency regularization prevents GNNs from overfitting to noisy labels via mimicry loss in both the inter-view and intra-view perspectives. To leverage such supervisions, we divide labels into clean and noisy types, rectify inaccurate labels, and further generate pseudo-labels on unlabeled nodes. Supervision for nodes with different types of labels is then chosen adaptively. This enables sufficient learning from clean labels while limiting the impact of noisy ones. We conduct extensive experiments to evaluate the effectiveness of our RTGNN framework, and the results validate its consistent superior performance over state-of-the-art methods with two types of label noises and various noise rates.

CCS CONCEPTS

• Information systems → Information retrieval.

*Corresponding Author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WSDM '23, February 27–March 3, 2023, Singapore, Singapore.

© 2023 Association for Computing Machinery.

ACM ISBN 978-1-4503-9407-9/23/02...\$15.00

<https://doi.org/10.1145/XXXXXX.XXXXXX>

KEYWORDS

GNNs, semi-supervised learning, label noise, robustness

ACM Reference Format:

Siyi Qian, Haochao Ying, Renjun Hu, Jingbo Zhou, Jintai Chen, Danny Z. Chen, and Jian Wu. 2023. Robust Training of Graph Neural Networks via Noise Governance. In *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining (WSDM '23)*, February 27–March 3, 2023, Singapore, Singapore. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/XXXXXX.XXXXXX>

1 INTRODUCTION

In real-world applications, a set of objects and their relationships can often be represented naturally as a graph. The graph data structure is widely employed in various domains such as biology, transportation, and social science. In the past decade, Graph Neural Networks (GNNs) have shown promising capacity in modeling graph data [13, 20, 38]. Typically, GNNs adopt a message passing and aggregation procedure to effectively propagate information via the graph structure. This mechanism makes GNNs very suitable for semi-supervised graph learning (e.g., node classification [1, 32]).

While GNNs are generally effective, most of the existing approaches assume that labels are sufficient and clean. However, in practice, node labels can be both scarce and noisy. For example, consider a graph from social media whose node labels are contributed by users. It is often the case that only a small fraction of users would participate in label generation, and intentionally or for other reasons, some of the labels do not reflect the truth. Another example is crowd-sourcing node labels, e.g., fake news annotation and medical knowledge graph annotation. It is easy to see that the annotation process is labor-intensive and expensive, and almost inevitably, label errors are introduced due to subjective judgment. Such cases will lead to graphs with scarce and noisy node labels. It has already been observed that noisy labels could pose a severe threat to the generalization performance of deep learning models [2, 42]. Therefore, developing label-efficient and noise-resistant GNNs is an important and challenging problem.

In the literature, robust deep learning in the presence of noisy labels has been explored mainly in applications with non-graph data such as images. Several strategies have been developed to combat label noises, e.g., sample selection [14, 16, 24], robust loss functions [10, 34, 44], and loss correction [27, 29]. However, simply incorporating these approaches into GNNs will be insufficient

(we will empirically show this in experiments). We observe that one unique characteristic of this problem with GNNs is that the scarcity of labels causes difficulties for nodes to receive sufficient supervision from labeled neighbors. Meanwhile, nodes in graphs are directly influenced by potential noises through message passing. Failing to balance the two would either lead to massive erroneous supervision from noisy labels or end up with insufficient learning.

Hence, a main challenge to solving this problem is how to *effectively leverage supervision of clean labels while limiting the impact of noisy ones*. Previous studies mainly focused on leveraging supervision of labels, but neglected how to limit the impact of noisy labels. For example, recently, NRGNN [6] investigated a robust GNN with noisy and sparse labels. It proposed to link unlabeled nodes with labeled nodes and further mine accurate pseudo-labels to provide more supervision. However, NRGNN mainly emphasized leveraging supervision of labels, which mixed up clean and noisy labels in its learning process. Whereas, as explained above, explicitly governing noises (i.e., limiting the impact of noisy labels) is necessary in order to further boost the robustness of GNNs.

To this end, in this paper, we propose a novel GNN framework called RTGNN (Robust Training of Graph Neural Networks via Noise Governance) that is capable of conducting robust learning with scarce and noisy node labels. As a distinguished design of our model, we develop a fine-grained noise governance strategy in semi-supervised learning. Due to the fact that deep neural networks (DNNs) tend to prioritize learning simple patterns first and may overfit to noises [2], we first propose an effective and scalable loss-based label division method to identify potentially noisy labels. Further, besides the commonly-used label supervision, we introduce self-reinforcement and consistency regularization as supplemental supervisions. The self-reinforcement supervision is inspired by the memorization effects of DNNs and aims to correct noisy labels. Consistency regularization includes an inter-view regularization based on ensembled classifiers capable of filtering different errors, and an intra-view regularization to explore the local homogeneity of the graph [39].

Specifically, we propose to train a pair of peer GNNs to enforce adequate supervision as well as govern noise labels. First, we augment the raw graph by linking labeled and unlabeled nodes to facilitate efficient message passing, following NRGNN [6]. Then, in each epoch, we progressively perform the main division of clean and noisy candidate set by the small-loss criterion [14] using a mutual loss. Next, a subset of confident nodes which predict a different class from their labels in the noisy candidate set is reinforced by training with their own prediction. Similarly, we extend the label set by adding those confident and consistent unlabeled nodes to the training set. Finally, we apply inter-view regularization to help two classifiers cooperatively mimic each other's soft targets and intra-view regularization to enable nodes to learn from their neighbors.

In summary, the main contributions of our work are as follows.

- We investigate the robust training problem of GNNs from the noise governance perspective, which has been under-explored in previous studies.
- We develop a novel RTGNN model which governs label noises explicitly with self-reinforcement and consistency

regularization. RTGNN also enables fine-grained learning on relatively clean, potentially noisy, and pseudo labels. This allows effectively leveraging supervision information while limiting the impact of label noises.

- We conduct extensive experiments to evaluate the effectiveness of our new approach, and the results validate the consistent superior performance of RTGNN over state-of-the-art methods with two types of noises and various noise rates.

2 RELATED WORK

2.1 Robust Deep Learning

It has been reported in [15, 22] that DNNs are prone to over-parameterize. As a result, they could overfit to and even memorize noisy labels, resulting in degraded generalization and poor performance [2, 42]. In the last few years, robust deep learning dealing with label noises is gaining more attention, especially in applications with non-graph data such as images [10, 14, 17, 24, 29, 44]. Such work can be roughly divided into two categories: sample-centric and loss-centric.

Sample-centric approaches seek to enhance the robustness of deep models by selecting clean samples for training. Malach *et al.* [24] proposed an update-by-disagreement strategy for sample selection, and trained two predictors simultaneously and performed parameter updates only on samples in disagreement areas of the two predictors. Inspired by the memorization effects of DNNs, Han *et al.* [14] developed a Co-teaching paradigm that also trains two networks such that each network samples its small-loss instances to update the parameters of its peer network. The small-loss Co-teaching paradigm has been further extended by bridging the update-by-disagreement strategy as Co-teaching+ [41] and the Co-regularization max-agreement principle as JoCoR [35]. Recently, more metrics have been explored for sample selection, e.g., loss at different stages of model fitness [16], network output at different training epochs [26], margins to decision boundaries [30], sample uncertainty [37], and divergence between different classifiers [40].

In contrast, loss-centric approaches adopt noise-robust loss functions or perform direct loss correction to combat label noises. One type of such methods aims to design loss functions that are robust to noise. Ghosh *et al.* [10] showed that the mean absolute error (MAE) loss is inherently robust to label noises. This idea was further extended by GCE [44] which can be viewed as a generalization of MAE and cross entropy (CE). SCE [34] boosts CE symmetrically with a noise-robust term called reverse cross entropy (RCE). The complementary loss function method [33] utilizes both CE loss and robust loss to balance learning sufficiency and robustness. APL [23] classifies existing loss functions into active and passive based on their optimization behavior; it leverages active normalized loss for robust training and complements training with passive loss to address the under-fitting problem. The second type of methods is direct loss correction. Goldberger *et al.* [11] adopted a noise adaptation layer to model and correct the noise distribution. The F-correction method [29] proposes forward and backward correction which rectifies loss based on the noise transition matrix.

While the above methods achieved good performance in supervised training on non-graph data, our work explores a new perspective that focuses on semi-supervised learning with noisy

labels on graphs and designs a new noise governance strategy to alleviate noise impact.

2.2 Training GNNs with Noisy Labels

GNNs have achieved state-of-the-art performance on several semi-supervised benchmarking tasks, e.g., link prediction and node classification [13, 20, 31, 38]. GNNs broaden the convolution operation to graph data. They learn a graph node’s representation by aggregating the features of its neighbors. GNNs can mainly be divided into spectral-based [3, 8, 20] and spatial-based [13, 31, 38] ones. While it has been shown that deep models are vulnerable to label noises, robust training of GNNs is generally under-explored, and only a few studies investigated the robustness of GNNs with noisy labels [6, 27]. Among them, D-GNN [27] shows that GCNs could be very vulnerable to label noises, and exploits backward loss correction [29] to boost performance. However, one-step loss correction is insufficient to deal with challenging scenarios, e.g., the one with noisy and scarce labels that we consider. Close to our work is NRGNN [6], which learns a robust GNN with noisy and sparse labels; it proposes to link unlabeled nodes with labeled ones and further mine accurate pseudo-labels to provide more supervision. However, noisy labels are not explicitly governed by NRGNN, thus resulting in error propagation during message passing to some extent. Different from NRGNN, we distinguish noisy labels from clean ones and apply additional supervision and label correction to reduce the impact of label noises.

3 PRELIMINARIES

Notation and problem formulation. We represent a graph $\mathcal{G} = (\mathcal{V}, \mathbf{A}, \mathbf{X})$ as a triple: a vertex set \mathcal{V} , an adjacency matrix \mathbf{A} , and a node feature matrix \mathbf{X} . Without loss of generality, let the vertices of \mathcal{G} be consecutively labeled from 1 to $|\mathcal{V}|$, i.e., $\mathcal{V} = \{1, 2, \dots, |\mathcal{V}|\}$. We consider undirected and unweighted graphs. Thus, the adjacency matrix $\mathbf{A} \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{V}|}$ is binary: $\mathbf{A}_{ij} = 1$ if there is an edge between nodes i and j , and $\mathbf{A}_{ij} = 0$ otherwise. Each node is assigned with a d -dimensional feature vector. The node feature matrix $\mathbf{X} \in \mathbb{R}^{|\mathcal{V}| \times d}$ arranges these feature vectors vertically such that the i -th row of \mathbf{X} , denoted as $\mathbf{X}^i \in \mathbb{R}^d$, is the feature vector of node i .

We focus on robust training of GNNs for semi-supervised node classification. It assumes that the nodes of a subset $\mathcal{V}_L \subset \mathcal{V}$ are assigned with labels in advance. Let $C > 1$ denote the number of node classes. Each node label $Y \in \{1, \dots, C\}$ can be alternatively represented as a one-hot vector $\mathbf{y} \in \{0, 1\}^C$. Recall that node labels can be scarce and noisy in practice. In other words, (1) $|\mathcal{V}_L|$ is significantly smaller than $|\mathcal{V}|$, and (2) \mathbf{y}^i might be incorrect for some of the nodes $i \in \mathcal{V}_L$. To tackle the node classification task, GNNs first compute a logit vector $\mathbf{o}_\theta^i \in \mathbb{R}^C$ for each node i , where θ is the model parameters. The probability of node i belonging to class j is then derived as:

$$\mathbf{p}_\theta^{i,j} = \exp(\mathbf{o}_\theta^{i,j}) / \sum_{k=1}^C \exp(\mathbf{o}_\theta^{i,k}), \quad (1)$$

where $\mathbf{o}_\theta^{i,k}$ denotes the k -th entry of \mathbf{o}_θ^i .

Two types of label noises. We consider two types of categorical label noises: *uniform* and *pair*. These two noise types are widely

adopted in the literature [6, 14, 24]. Let ϵ denote the noise rate of node labels. For uniform noise, the labels have a probability of $\epsilon/(C-1)$ to be uniformly flipped to other classes. For pair noise, the labels probably make mistakes only within very similar class [14], i.e., the labels have a probability of ϵ to be flipped to the paired class. *In this work, we assume that the majority of labeled nodes have correct labels, i.e., the noise rate $\epsilon < 0.5$.*

4 METHODOLOGY

4.1 Overview

Our new RTGNN framework for Robust Training of GNNs with scarce and noisy node labels seeks to explicitly govern label noise to enable sufficient learning from clean labels while limiting the impact of noisy ones. An overview of RTGNN is given in Fig. 1.

RTGNN first augments the input graph by learning an edge predictor to infer potential links between labeled and unlabeled nodes (Sec. 4.2). These added links lead to more efficient message passing, thus mitigating the label scarcity issue. The augmented graph is then fed to a pair of peer GCNs for explicit noise governance (Sec. 4.3). Based on the prediction of the two peer GCNs, nodes are classified into different categories, enabling nodes to choose supervision strategies adaptively. More specifically, labeled nodes are divided into clean candidate set \mathcal{V}_{cl} and noisy candidate set \mathcal{V}_{ns} based on the small-loss criterion [14] (Sec. 4.3.1). The clean nodes are mainly supervised by their assigned labels. RTGNN further identifies a subset $\mathcal{V}_{sr} \subseteq \mathcal{V}_{ns}$ of noisy label nodes whose predictions are confident but different from their assigned labels. It then reinforces the training of nodes in \mathcal{V}_{sr} with their own predicted labels to alleviate noise propagation (Sec. 4.3.2).

For the remaining nodes in $\mathcal{V}_{ns} \setminus \mathcal{V}_{sr}$, label-based training is performed with a down-weighted loss. Similar to the self-reinforce case, RTGNN further generates pseudo-labels for those unlabeled nodes whose predictions are confident (Sec. 4.3.3) in order to facilitate sufficient learning. Finally, it introduces the inter-view and intra-view regularizations to further prevent the model from overfitting to noise (Sec. 4.3.4). We will summarize the different categories of nodes and their supervisions in Sec. 4.4.

4.2 Graph Augmentation

Latent graph learning has shown to be beneficial in a number of graph learning tasks, owing to its flexibility to form graph structures optimized for specific tasks [5, 9, 19, 21, 28, 45]. Accordingly, we also utilize latent graph learning to densify the input graph as augmentation, to promote supervision propagation and alleviate label scarcity. Inspired by [6], we choose to infer potential links between labeled and unlabeled nodes. The reasons are two-fold: 1) labeled nodes could provide direct data supervision to similar unlabeled ones, and 2) unlabeled nodes with pseudo-labels (Sec. 4.3.3) could facilitate model supervision on labeled ones as a supplement.

To keep the extra computational cost of graph augmentation affordable, we decouple the process into candidate generation and link inference. For candidate generation, we retrieve the top- K nearest labeled/unlabeled nodes for each unlabeled/labeled one and the distance is evaluated in the node feature space [7, 9]. The generation step can be conducted offline in advance and efficient approximation methods [12, 46] are available for large graphs.

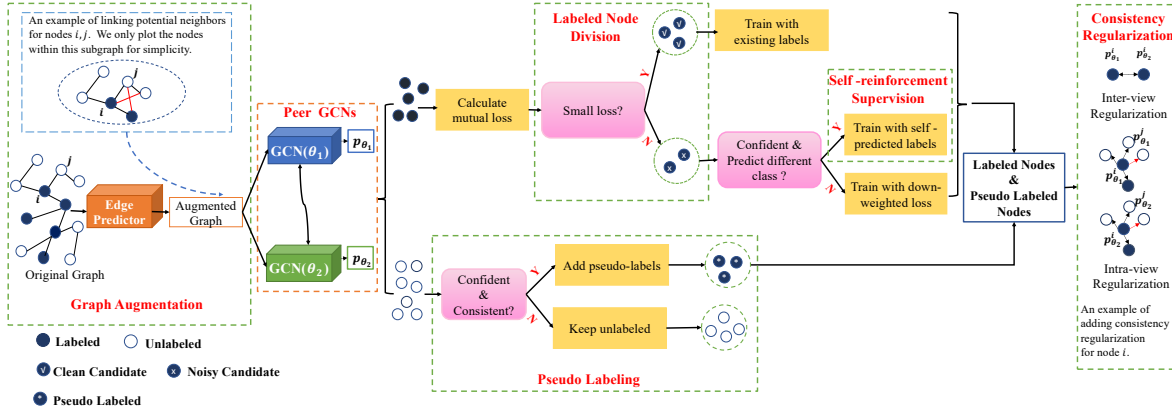


Figure 1: The overall architecture of RTGNN. The red lines denote the added links. RTGNN particularly leverages a pair of peer GCNs for noise governance.

Next, we train an encoder-decoder [21] edge predictor to infer links. A GCN [20] is utilized as the encoder to project nodes into a latent space based on both node feature and structure information:

$$\mathbf{Z} = \text{GCN}(\mathbf{X}, \mathbf{A}). \quad (2)$$

We then adopt an inner-product decoder for link inference. More specifically, the *predicted* edge weight w_{ij} between nodes i and j is estimated as the non-negative cosine similarity between their node representations \mathbf{z}^i and \mathbf{z}^j learned by the encoder:

$$w_{ij} = \max\left(\frac{\mathbf{z}^i \cdot \mathbf{z}^j}{\|\mathbf{z}^i\| \|\mathbf{z}^j\|}, 0\right). \quad (3)$$

We train the edge predictor with a negative-sampling based reconstruction objective:

$$\mathcal{L}_{rec} = \sum_i \left(\sum_{j: A_{ij} > 0} (w_{ij} - A_{ij})^2 + N_{neg} \cdot \mathbb{E}_{j' \sim P_n} (w_{ij'} - A_{ij'})^2 \right), \quad (4)$$

where N_{neg} is the number of negative samples for each node and P_n is the distribution of negative samples (e.g., uniform). We employ negative sampling to improve computational efficiency and avoid bias towards negative node pairs. The adjacent matrix $\hat{\mathbf{A}}$ of the augmented graph is constructed as:

$$\hat{A}_{ij} = \begin{cases} 1 & \text{if } A_{ij} = 1; \\ w_{ij} & \text{if } A_{ij} = 0, j \in \text{top}K(i), \text{ and } w_{ij} > \tau; \\ 0 & \text{otherwise.} \end{cases} \quad (5)$$

Note that $\text{top}K(i)$ includes the retrieved nearest neighbor candidates for node i and the threshold τ filters those unreliable ones.

Following [4, 6], we train an end-to-end model by jointly learning node classification and graph reconstruction. The training process will be described in Sec. 4.4.

4.3 Robust Training via Noise Governance

The core of RTGNN lies in explicit label noise governance. To achieve this, RTGNN feeds the augmented graph to a pair of peer GCNs that share the same architecture but have different parameters, θ_1 and θ_2 . Based on the predictions of these two GCNs, says \mathbf{p}_{θ_1} and \mathbf{p}_{θ_2} , RTGNN classifies the nodes into different categories and further exploits different supervision strategies accordingly. As

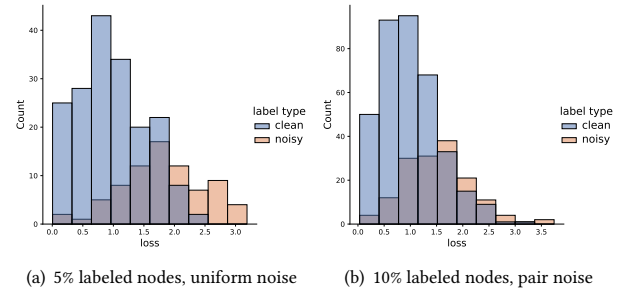


Figure 2: Loss distribution of the 40th epoch when training a GCN on BlogCatalog [36] with a noise rate $\epsilon = 0.3$.

such, RTGNN is able to effectively exploit the supervision of clean labels while limiting the impact of noisy ones.

4.3.1 Labeled node division. We first divide labeled nodes of \mathcal{V}_L into clean and noisy candidate sets, \mathcal{V}_{cl} and \mathcal{V}_{ns} . Previous studies have observed that DNNs follow the pattern of learning clean labels first and then gradually overfitting to noisy labels [2]. This will lead to different loss distributions for clean and noisy nodes. Fig. 2 also empirically verifies this phenomenon when training a GCN on graph data: The loss values of clean nodes are generally smaller than those of noisy nodes in the early stage. Inspired by the above observation, we adopt the small-loss criterion [14] for labeled node division. Also, here we leverage the peer GCNs as different classifiers possess different abilities to exclude noises because they form diverse decision boundaries [14, 41]. We define the mutual loss derived from cross entropy as the loss metric for a node i , as:

$$\mathcal{L}_{mutual}^i = -\mathbf{y}^i \log(\mathbf{p}_{\theta_1}^i) - \mathbf{y}^i \log(\mathbf{p}_{\theta_2}^i) = -\mathbf{y}^i \log(\mathbf{p}_{\theta_1}^i \cdot \mathbf{p}_{\theta_2}^i). \quad (6)$$

The mutual loss measures prediction confidence, i.e., \mathcal{L}_{mutual}^i has a low value if both the GCN classifiers correctly predict the class of node i . As clean labels are easier to learn, we expect that the mutual loss values for clean label nodes are smaller than noisy ones.

Let T_{max} denote the number of training epochs. In the t -th epoch, clean and noisy candidate sets are formally divided as follows:

$$\begin{aligned} th_{epoch}^t &= \text{Percentile}(\mathcal{L}_{mutual}^i, 100 - \frac{50t}{T_{max}}), \\ th_{avg}^t &= \text{Average}(\mathcal{L}_{mutual}^i), \\ \mathcal{V}_{cl} &= \{i \mid \mathcal{L}_{mutual}^i < \max(th_{epoch}^t, th_{avg}^t)\}, \\ \mathcal{V}_{ns} &= \mathcal{V}_L \setminus \mathcal{V}_{cl}. \end{aligned} \quad (7)$$

Notes: (1) $\text{Percentile}(\mathcal{L}, p)$ is the value below which $p\%$ of the loss values in \mathcal{L} fall, (2) thresholds th_{epoch}^t and th_{avg}^t are the upper bounds of loss values for clean label nodes, (3) th_{epoch}^t controls that the proportion of clean label nodes decreases with the increment of the training epoch t , (4) at least half of labeled nodes are classified as clean given a noise rate $\epsilon < 0.5$, and (5) the threshold th_{avg}^t further ensures that small-loss nodes are clean regardless of their relative ranks.

The training on nodes in \mathcal{V}_{cl} is supervised by their assigned labels, with \mathcal{L}_{mutual}^i being the classification loss. On the other hand, we leverage different supervisions for nodes in \mathcal{V}_{ns} , which will be discussed in the following subsections. The above node division strategy has two advantages. First, it does not require hyper-parameters, thus avoiding elaborate parameter tuning. Second, it strikes a balance between sufficient training in the early stage and noise-resistant training in the later stage. At the early stage of training (e.g., the first few epochs), the model is underfitting and has not yet learned reliable patterns. By choosing a conservative (i.e., high) th_{epoch}^t , the majority of labeled nodes are classified as clean to guide model supervision. This can also be viewed as warm-up training. Besides, as mentioned before, the memorization effect of neural networks enables one to learn from clean labels first, despite the existence of noise. Thus, choosing a conservative th_{epoch}^t would not hurt the model performance. As the training continues, we progressively exclude nodes with large loss from \mathcal{V}_{cl} , which prevents our model from overfitting to noise in the later stage.

4.3.2 Self-reinforcement supervision. As the labels of nodes in \mathcal{V}_{ns} are prone to be noisy, RTGNN turns to other supervisions instead. Note that properly trained deep models have the ability to predict the correct labels on their own. Thus, we give a chance for nodes in \mathcal{V}_{ns} to reinforce the training with their own predicted labels. We identify a subset $\mathcal{V}_{sr} \subseteq \mathcal{V}_{ns}$ of nodes on which the predictions of the two peer GCNs are confident but different from their labels. Formally, \mathcal{V}_{sr} includes the nodes that satisfy:

$$\begin{aligned} Z^i &= \arg \max_c \mathbf{p}_{\theta_1}^{i,c} = \arg \max_c \mathbf{p}_{\theta_2}^{i,c} \neq Y^i, \\ \sqrt{\mathbf{p}_{\theta_1}^{i,Z^i} \cdot \mathbf{p}_{\theta_2}^{i,Z^i}} &> 1 - \frac{(C-1)t}{C \cdot T_{max}}, \end{aligned} \quad (8)$$

where $Y^i \in \{1, \dots, C\}$ denotes the label for node i . The threshold $1 - \frac{(C-1)t}{C \cdot T_{max}}$ linearly decreases with the increment of t and is finally reduced to $1/C$. In other words, we gradually allow more nodes to self-reinforce the training.

In each epoch, we update \mathcal{V}_{sr} and further compute an adaptive training weight for each node $i \in \mathcal{V}_{sr}$:

$$\begin{aligned} \mu(i) &= (\sqrt{\mathbf{p}_{\theta_1}^{i,Z^i} \cdot \mathbf{p}_{\theta_2}^{i,Z^i}})^{1 - \frac{t}{T_{max}}}, \\ \mathcal{L}_{sr}^i &= -\mu(i) \mathbf{z}^i \log(\mathbf{p}_{\theta_1}^i \cdot \mathbf{p}_{\theta_2}^i), \end{aligned} \quad (9)$$

where $\mathbf{z}^i \in \{0, 1\}^C$ is the corresponding one-hot representation of the predicted label Z^i , and $\mu(i)$ serves as an adaptive learning rate. Notes: (1) $\mu(i)$ relies on the node-wise prediction confidence. (2) The scaling exponent $(1 - t/T_{max})$ gradually decreases to zero to enlarge $\mu(i)$ in the later training stage.

4.3.3 Pseudo labeling. Learning from pseudo-labels may further ease the training with scarce and noisy labels [6]. Similarly, we select a subset $\mathcal{V}_{pse} \subset \mathcal{V} \setminus \mathcal{V}_L$ of unlabeled nodes with confident and consistent predictions to generate pseudo-labels. Formally, \mathcal{V}_{pse} includes the nodes that satisfy:

$$\tilde{Z}^i = \arg \max_c \mathbf{p}_{\theta_1}^{i,c} = \arg \max_c \mathbf{p}_{\theta_2}^{i,c}, \sqrt{\mathbf{p}_{\theta_1}^{i,\tilde{Z}^i} \cdot \mathbf{p}_{\theta_2}^{i,\tilde{Z}^i}} > th_{pse}, \quad (10)$$

where th_{pse} refers to the confidence threshold. The training on nodes in \mathcal{V}_{pse} is then supervised by their pseudo-labels.

4.3.4 Consistency regularization. For nodes in $\mathcal{V}_L \cup \mathcal{V}_{pse}$, we further add some regularization to prevent overfitting. Inspired by mutual learning [43], we let the two peer GCNs mimic each other's prediction. In this way, they teach and learn from each other. Here we use D_{KL} (Kullback-Liebler divergence) as the mimicry loss:

$$\mathcal{L}_{inter}^i = D_{KL}(\mathbf{p}_{\theta_1}^i \parallel \mathbf{p}_{\theta_2}^i) + D_{KL}(\mathbf{p}_{\theta_2}^i \parallel \mathbf{p}_{\theta_1}^i), \quad (11)$$

Moreover, we follow the local consistency (a.k.a. homophily [25]) assumption that linked nodes tend to belong to the same classes. This leads to a regularization term that enforces nodes to mimic neighbors' predictions within the same GCN:

$$\mathcal{L}_{intra}^i = \sum_j \frac{\hat{A}_{ij}}{\sum_k \hat{A}_{ik}} (D_{KL}(\mathbf{p}_{\theta_1}^j \parallel \mathbf{p}_{\theta_1}^i) + D_{KL}(\mathbf{p}_{\theta_2}^j \parallel \mathbf{p}_{\theta_2}^i)). \quad (12)$$

The final regularization loss is composed as:

$$\mathcal{L}_{reg}^i = \mathcal{L}_{inter}^i + \mathcal{L}_{intra}^i. \quad (13)$$

4.4 The Overall Training Loss

Combining all together, the loss of labeled nodes can be unified as:

$$\mathcal{L}_{labeled} = -\frac{1}{|\mathcal{V}_L|} \sum_{i \in \mathcal{V}_L} (\zeta(i) \hat{\mathbf{y}}^i \log(\mathbf{p}_{\theta_1}^i \cdot \mathbf{p}_{\theta_2}^i) + \lambda \mathcal{L}_{reg}^i), \quad (14)$$

where

$$\begin{cases} \zeta(i) = 1, \hat{\mathbf{y}}^i = \mathbf{y}^i & \text{if } i \in \mathcal{V}_{cl}; \\ \zeta(i) = \mu(i), \hat{\mathbf{y}}^i = \mathbf{z}^i & \text{if } i \in \mathcal{V}_{sr}; \\ \zeta(i) = \gamma \in (0, 1), \hat{\mathbf{y}}^i = \mathbf{y}^i & \text{if } i \in \mathcal{V}_{ns} \setminus \mathcal{V}_{sr}. \end{cases}$$

Note that γ is a down-weighted factor for nodes with noisy labels and unconfident predictions.

Similarly, the training loss for nodes with pseudo-labels is:

$$\mathcal{L}_{pse} = -\frac{1}{|\mathcal{V}_{pse}|} \sum_{i \in \mathcal{V}_{pse}} (\tilde{\mathbf{z}}^i \log(\mathbf{p}_{\theta_1}^i \cdot \mathbf{p}_{\theta_2}^i) + \lambda \mathcal{L}_{reg}^i), \quad (15)$$

Table 1: Dataset statistics (the data is cited from [18, 36]).

Dataset	Nodes	Edges	Features	Classes
Cora	2,485	5,069	1,433	7
Citeseer	2,110	3,668	3,703	6
BlogCatalog	5,196	171,743	8,189	6

where $\tilde{\mathbf{z}}^i \in \{0, 1\}^C$ is the corresponding one-hot representation of the predicted pseudo-label \tilde{Z}^i (Eq. (10)).

Combining the graph reconstruction loss (Eq. (4)) given in Sec. 4.2, the total training loss is:

$$\mathcal{L}_{total} = \mathcal{L}_{labeled} + \mathcal{L}_{pse} + \alpha \cdot \mathcal{L}_{rec}. \quad (16)$$

5 EXPERIMENTS

5.1 Experiment Setup

Datasets. We conduct experiments on two citation datasets [18] (Cora and Citeseer) and a social network dataset BlogCatalog [36] to evaluate the performance of RTGNN. The dataset statistics are summarized in Table 1. To study the robustness of GNNs with scarce and noisy labels, each dataset is randomly split into 5%, 15%, and 80% for training, validation, and test, respectively. Moreover, following [6], we further randomly corrupt a fraction of labels, says ϵ , in the training and validation sets. Recall in Sec. 3 that we consider two types of noises. Each dataset thus has two corrupted versions corresponding to uniform and pair noises, respectively.

Implementation details. We adopt a 2-layer GCN [20] with 128 hidden units as the backbone GNN model. For graph augmentation, we use a GCN with 64 hidden units as the encoder of edge predictor, fix $N_{neg} = 100$ for negative sampling, set threshold $\tau = 0.05$ on all datasets. The regularization loss weight $\lambda = 0.1$ for all datasets.

Moreover, we search the down-weighted factor γ of noisy samples in $\{0.01, 0.1\}$, the K for nearest neighbor candidate (the $topK$ operation in Eq. (5)) in $\{25, 50, 100\}$, the pseudo label threshold th_{pse} in $\{0.7, 0.8, 0.9, 0.95\}$, and the reconstruction loss weight α in $\{0.03, 0.1, 0.3, 1\}$ based on the validation performance. Finally, we train our model for a total of 200 epochs with a learning rate of 0.001 and a weight decay of $5e-4$. To avoid over-fitting, we also apply dropout with a dropout rate of 0.5. In the cases when quantitative measurements are reported, the test was repeated over 5 times and the average is reported. **Notice that both GCNs could be used for inference, and in the experiments, we use the first one.**

5.2 Performance Comparison with Baselines

We first compare our RTGNN with the following state-of-the-art robust learning baselines to evaluate the overall performance.

- **GCN [20]** is a popular GNN model based on first-order approximate spectral convolution.
- **Co-teaching [14]** trains paired peer networks such that each network excludes a proportion (based on the given noise rate) of large-loss nodes and uses the rest to update the parameters of the peer network.
- **JoCoR [35]** improves Co-teaching by jointly training and discarding a proportion of large-loss nodes. It also adopts the paired network architecture and assumes a given noise rate.
- **SCE [34]** enhances noise robustness by adding a noise-tolerant reverse cross entropy term to the cross entropy loss.

- **APL [23]** finds that normalization helps improve noise robustness of loss functions. It then categorizes robust loss functions into active and passive and combines them for training. In particular, it leverages a combination of normalized cross entropy and reverse cross entropy.
- **D-GNN [27]** introduces backward loss correction [29] to GNNs, which multiplies the cross entropy loss with the inverse of the noise matrix to alleviate loss distribution bias.
- **NRGNN [6]** adopts two asymmetric GCNs for edge prediction and pseudo-label mining. It emphasizes on extracting and broadcasting more potentially correct supervision on graphs while neglecting proactive governance of noises.

We use the official implementations of these baselines. For a fair comparison, all the methods adopt a 2-layer GCN with 128 hidden units as the backbone model. We evaluate the node classification accuracy of all methods on three datasets with two types of noises and varying noise rates. The results are reported in Table 2 and we observe the following from these results.


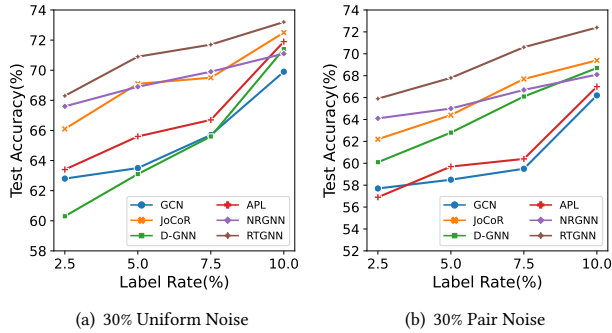
- Compared with GCN [20], loss-centric methods SCE, APL, and D-GNN perform even worse in some scenarios (e.g., SCE on Cora with 40% pair noise, APL on Cora with 40% uniform noise, and D-GNN on Citeseer with 40% pair noise). It indicates that leveraging a single robust loss function or loss correction method is insufficient. Differently, RTGNN explicitly divides nodes into several categories and adopts different training strategies accordingly.
- Compared with GCN, both Co-teaching and JoCoR yield stable and superior performance. This verifies that co-training two networks can mutually suppress the effects of noise, inspiring us to borrow the peer network architecture in our model. It also validates that performing explicit governance of clean and noisy labels is beneficial.
- JoCoR attains an overall better performance than Co-teaching by jointly training two networks together. This motivates us to fuse the training of two peer networks. Moreover, our RTGNN works even better, because we further leverage the model’s own ability to rectify the noisy labels. Besides, we complement the training with confident and consistent pseudo labels.
- NRGNN links unlabeled nodes with similar labeled nodes and adds pseudo-labels to mitigate the negative effect of noise. This strategy helps to bring effective supervision to unlabeled nodes. However, it might not perform well under higher noise rates since a large number of erroneous supervisions are propagated to a wide scope. To address this issue, we propose explicit noise governance to identify potentially noisy labels and alleviate their impacts. We consider our labeled node division as a primary module in RTGNN, with other types of supervisions being introduced to further boost the performance. **Overall, RTGNN consistently performs better than all baselines on three datasets with both types of noises and varying noise rates.**

5.3 Impacts of Training Label Rates

We next examine how our RTGNN performs with different fractions of training labels. For this purpose, we split the datasets into $x\%$, $(20 - x)\%$, and 80% for training, validation, and test, respectively. Since we focus on training GNNs with sparse and noisy labels, we vary x from 2.5 to 10 and evaluate the test node classification

Table 2: Comparison of node classification performance (test accuracy, mean \pm std, averaged by 5 runs) by different models. The best results are marked in bold, and the best baseline results are underlined.

Dataset	Model	Uniform Noise				Pair Noise			
		10%	20%	30%	40%	10%	20%	30%	40%
Cora	GCN	79.0 \pm 0.3	76.6 \pm 0.5	68.8 \pm 0.9	60.1 \pm 3.0	79.8 \pm 0.4	75.0 \pm 1.4	67.0 \pm 1.0	59.3 \pm 0.6
	Co-teaching	79.2 \pm 0.4	77.0 \pm 0.3	69.5 \pm 0.7	62.7 \pm 1.8	80.6 \pm 0.5	76.7 \pm 0.8	67.7 \pm 0.7	63.8 \pm 0.8
	JoCoR	79.3 \pm 0.2	76.9 \pm 0.3	72.9 \pm 1.4	68.2 \pm 2.5	80.7 \pm 0.3	77.6 \pm 0.3	69.2 \pm 1.4	63.8 \pm 0.9
	SCE	78.9 \pm 0.4	77.0 \pm 0.8	69.8 \pm 0.9	60.0 \pm 3.4	80.5 \pm 0.2	77.4 \pm 1.0	67.8 \pm 2.4	57.8 \pm 1.3
	APL	79.4 \pm 0.3	74.4 \pm 0.6	69.2 \pm 1.4	59.5 \pm 3.2	80.0 \pm 0.5	75.2 \pm 2.3	67.4 \pm 0.9	59.8 \pm 2.2
	D-GNN	79.6 \pm 0.1	74.7 \pm 0.6	71.4 \pm 0.5	61.0 \pm 0.6	78.5 \pm 0.9	76.4 \pm 0.4	68.9 \pm 1.2	59.1 \pm 1.6
	NRGNN	81.7 \pm 0.4	80.4 \pm 1.2	76.7 \pm 0.7	73.2 \pm 1.9	81.7 \pm 0.4	79.5 \pm 0.8	75.4 \pm 1.1	68.1 \pm 1.6
	RTGNN (ours)	82.6\pm0.7	80.8\pm0.2	79.9\pm1.0	78.1\pm1.1	82.9\pm0.3	80.9\pm0.6	77.4\pm0.6	72.1\pm1.4
Citeseer	GCN	69.8 \pm 0.7	69.4 \pm 0.4	63.5 \pm 1.2	55.0 \pm 1.0	69.2 \pm 0.5	67.0 \pm 0.8	58.5 \pm 2.1	57.5 \pm 1.6
	Co-teaching	72.1 \pm 0.3	69.6 \pm 1.0	69.0 \pm 1.6	56.9 \pm 1.4	72.2 \pm 0.5	69.3 \pm 1.3	60.1 \pm 1.7	57.6 \pm 2.4
	JoCoR	72.2 \pm 0.3	71.3 \pm 0.7	<u>69.1\pm0.3</u>	58.1 \pm 2.3	72.1 \pm 0.6	69.6 \pm 0.4	64.4 \pm 2.5	58.1 \pm 1.6
	SCE	70.2 \pm 0.4	69.4 \pm 0.7	64.9 \pm 1.6	56.2 \pm 1.4	71.1 \pm 0.2	67.8 \pm 0.8	59.8 \pm 0.6	57.9 \pm 1.0
	APL	70.5 \pm 0.6	69.6 \pm 0.6	65.6 \pm 1.6	56.7 \pm 2.0	71.3 \pm 0.4	68.1 \pm 0.5	59.7 \pm 0.6	57.1 \pm 1.5
	D-GNN	71.0 \pm 0.3	68.7 \pm 0.7	63.1 \pm 0.5	54.5 \pm 0.9	70.0 \pm 0.5	69.5 \pm 0.5	62.8 \pm 0.6	56.4 \pm 4.0
	NRGNN	<u>72.6\pm0.8</u>	<u>72.4\pm0.8</u>	68.9 \pm 1.1	<u>63.5\pm1.3</u>	<u>72.8\pm0.5</u>	<u>70.4\pm0.8</u>	<u>65.0\pm1.3</u>	<u>58.4\pm3.4</u>
	RTGNN (ours)	74.5\pm0.6	74.1\pm0.6	70.9\pm0.7	66.0\pm1.7	74.6\pm0.5	71.8\pm1.0	67.8\pm0.8	62.1\pm1.6
BlogCatalog	GCN	69.9 \pm 0.7	66.4 \pm 0.8	65.9 \pm 0.9	65.3 \pm 0.9	69.9 \pm 0.4	62.9 \pm 0.5	58.3 \pm 0.8	57.3 \pm 1.2
	Co-teaching	70.9 \pm 0.3	69.6 \pm 0.5	68.7 \pm 0.6	66.3 \pm 0.6	70.2 \pm 0.7	65.3 \pm 0.9	58.8 \pm 1.0	57.8 \pm 1.5
	JoCoR	70.9 \pm 0.4	69.6 \pm 0.4	69.3 \pm 1.1	66.3 \pm 1.1	70.5 \pm 0.7	66.3 \pm 1.1	59.5 \pm 1.1	58.3 \pm 2.2
	SCE	70.7 \pm 0.5	68.7 \pm 0.8	67.5 \pm 0.9	66.5 \pm 1.0	69.4 \pm 0.7	64.3 \pm 1.4	60.6 \pm 1.0	57.0 \pm 0.9
	APL	70.9 \pm 0.5	69.2 \pm 0.4	68.7 \pm 0.6	65.8 \pm 0.9	70.3 \pm 0.6	<u>68.4\pm0.8</u>	<u>61.8\pm1.1</u>	58.0 \pm 2.7
	D-GNN	70.7 \pm 0.3	67.9 \pm 0.6	67.5 \pm 0.5	65.8 \pm 0.7	70.1 \pm 0.5	67.1 \pm 1.1	61.5 \pm 1.3	58.4 \pm 1.6
	NRGNN	<u>71.1\pm0.4</u>	<u>70.3\pm1.0</u>	<u>69.5\pm0.9</u>	<u>67.0\pm1.3</u>	69.9 \pm 0.7	67.9 \pm 1.0	60.7 \pm 1.0	58.3 \pm 2.2
	RTGNN (ours)	71.2\pm0.3	70.9\pm0.4	70.7\pm0.4	68.0\pm1.1	71.1\pm0.4	70.3\pm0.8	65.2\pm1.0	63.6\pm1.5

**Figure 3: Impacts of training label rates on Cora.****Figure 4: Impacts of training label rates on Citeseer.**

accuracy. The results on Cora and Citeseer with 30% noise rate are

Table 3: Ablation study of graph augmentation. “w/ GA” means plugging the graph augmentation module.

Dataset	Variants	Uniform Noise		Pair Noise	
		30%	40%	30%	40%
Cite-seer	CE	63.5 \pm 1.2	55.0 \pm 1.0	58.5 \pm 2.1	57.5 \pm 1.6
	CE w/ GA	67.9\pm1.6	62.8\pm1.9	64.6\pm3.2	58.6\pm3.2
	SCE	64.9 \pm 1.6	56.2 \pm 1.4	59.8 \pm 0.6	57.9 \pm 1.0
	SCE w/ GA	69.2\pm1.7	63.9\pm1.0	64.8\pm1.4	60.4\pm3.0
	APL	65.6 \pm 1.6	56.7 \pm 2.0	59.7 \pm 0.6	57.1 \pm 1.5
	APL w/ GA	69.5\pm1.2	63.2\pm0.9	64.3\pm1.2	59.5\pm1.6
Blog-Catalog	CE	65.9 \pm 0.9	65.3 \pm 0.9	58.3 \pm 0.8	57.3 \pm 1.2
	CE w/ GA	68.0\pm0.4	65.9\pm1.9	59.8\pm0.7	58.1\pm1.7
	SCE	67.5 \pm 0.9	66.5 \pm 1.0	60.6 \pm 1.0	57.0 \pm 0.9
	SCE w/ GA	69.8\pm1.0	67.1\pm0.3	62.6\pm0.7	58.7\pm0.5
	APL	68.7 \pm 0.6	65.8 \pm 0.9	61.8 \pm 1.1	58.0 \pm 2.7
	APL w/ GA	69.6\pm1.0	67.3\pm0.6	62.6\pm1.0	59.1\pm0.5

reported in Fig. 3 and Fig. 4. Note that results are similar with other noise rates and we observe the following.

- Our proposed RTGNN consistently outperforms all baselines. Even under an extremely low label rate ($x = 2.5\%$), our model achieves favorable performances by leveraging the supervision of clean labels while restricting that of noisy ones.
- Compared using the low (2.5%) and high (10%) label rates, we observe that our model makes a bigger improvement upon the competitive NRGNN baseline with high (10%) label rates. As we mentioned before, without explicit noise governance, NRGNN mix-ups the propagation of clean and noisy labels in graphs, thus being severely affected by the increasing size of noisy labels.

5.4 Ablation Study

Table 4: Ablation study of noise governance. “w/o” denotes “without”.

Dataset	Variants	Uniform Noise		Pair Noise	
		30%	40%	30%	40%
Cite-seer	w/o LD, SR	69.5±0.8	64.0±1.3	66.2±0.7	59.7±1.3
	w/o SR	70.3±0.9	65.1±1.3	67.2±0.7	61.1±1.5
	w/o PL	70.5±1.2	65.7±1.5	67.3±1.7	61.4±1.8
	w/o CR	70.6±1.0	65.4±2.1	67.3±1.2	61.2±1.9
	All	70.9±0.7	66.0±1.7	67.8±0.8	62.1±1.6
Blog-Catalog	w/o LD, SR	69.2±0.5	67.1±0.6	61.4±0.5	58.5±1.2
	w/o SR	70.1±0.4	67.5±0.9	64.5±1.2	62.0±1.3
	w/o PL	70.4±0.4	67.5±1.3	64.6±1.1	62.9±2.0
	w/o CR	70.3±0.2	67.8±1.1	64.4±1.4	62.5±1.7
	All	70.7±0.4	68.0±1.1	65.2±1.0	63.6±1.5

5.4.1 Ablation study on graph augmentation. To validate the effect of the graph augmentation module, we plug it into a GCN. In addition to the cross entropy (CE) loss, various loss functions are also tested. As in Eq. (16), we train the augmented baselines by jointly learning node classification and graph reconstruction. Similarly, we fine-tune the loss weight α and K for the nearest neighbor candidate. We set the label rate as $x = 5\%$, and report the results in Table 3. We observe the following.

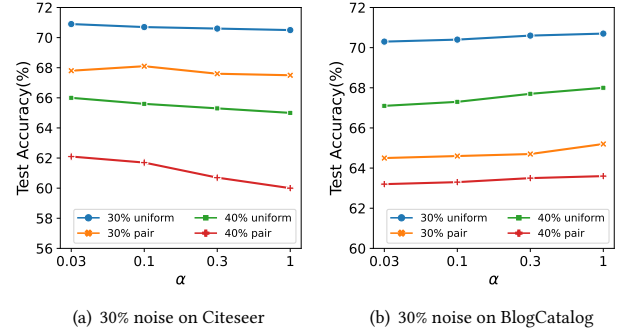
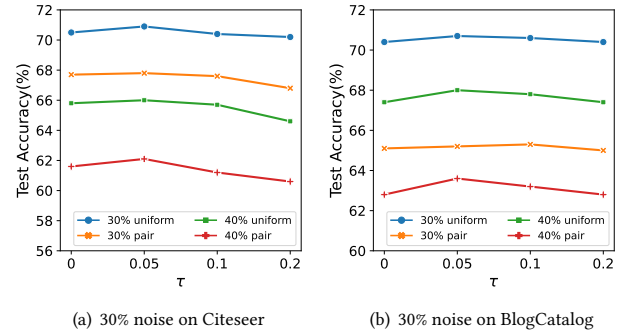

- The graph augmentation module boosts the training of GCN with scarce and noisy labels using various loss functions.
- A sparse graph (Citeseer) gains more benefits than a denser one (BlogCatalog). On a dense graph, abundant unlabeled nodes can receive sufficient but probably wrong supervision from labeled neighbors. Without an appropriate strategy to govern label noise, the chance for unlabeled nodes to receive correct supervision is not significantly increased, thus yielding limited performance improvement. This suggests that developing an effective noise governance strategy can largely help deal with graphs of different densities.

5.4.2 Ablation study on noise governance. We further investigate how each sub-module in our noise governance strategy contributes to the overall performance of RTGNN. For simplicity, we denote these sub-modules as follows.

- **LD:** Labeled node division (introduced in Sec. 4.3.1).
- **SR:** Self-reinforcement supervision (introduced in Sec. 4.3.2). It is an operation following labeled node division.
- **PL:** Pseudo labeling (introduced in Sec. 4.3.3).
- **CR:** Consistency regularization (introduced in Sec 4.3.4).

We set the label rate as $x = 5\%$, and present the results in Table 4. From these results, we observe the following.

- The RTGNN model without our labeled node division (LD) performs the worst. In particular, under 40% pair noise, it drops by 2.4% on Citeseer and by 5.1% on BlogCatalog. This indicates the importance of dividing labeled nodes into clean and noisy candidate sets and leveraging different training strategies for them. Meanwhile, the effectiveness of our designed labeled node division strategy is also validated.
- Self-reinforcement supervision (SR) further exploits the neural network’s ability for reinforcing itself with rectified labels, thus attaining more accurate supervision.
- Pseudo labeling (PL) enhances the performance because it serves as a kind of complementary supervision.

**Figure 5: Impact of α .****Figure 6: Impact of τ (averaged by 5 runs).****Figure 7: Impact of λ (averaged by 5 runs).**

- Consistency regularization (CR) further prevents nodes from overfitting to noise by ensembling two classifiers to exclude noise together and exploiting local homogeneity in the graph.

5.5 Sensitivity Analysis of Hyper-parameters

In this section, we conduct experiments to examine the impact of several important hyper-parameters in our model. The label rate is set as $x = 5\%$ and the noise rate is set as 30% and 40% as examples. We report the average results of 5 runs.

We consider the following key hyper-parameters. (1) In Fig. 5, we show the impact of the graph reconstruction loss weight α . On Citeseer, our method achieves the best results in **most** cases when $\alpha = 0.03$. On BlogCatalog, our method achieves the best results when $\alpha = 1$ for all the cases. (2) In Fig. 6, we show the impact of

the edge threshold τ . When $\tau = 0$, noisy edges may be introduced, and when τ becomes too large, some potential links would not be identified. On Citeseer, our method achieves the best results when $\tau = 0.05$ for all the cases. On BlogCatalog, our method achieves the best results in **most** cases when $\tau = 0.05$. (3) In Fig. 7, we show the impact of the consistency regularization weight λ . With the increase of λ , the model performance first increases and then drops. Our method achieves the best result when $\lambda = 0.2$ on Citeseer under 30% uniform and pair noise and $\lambda = 0.1$ on Citeseer under 40% uniform and pair noise. In general, $\lambda = 0.1$ gives good performances for various scenarios on Citeseer. Our method achieves the best results when $\lambda = 0.1$ on BlogCatalog for all the cases.

6 CONCLUSIONS

In this paper, we investigated how to train a robust GNN classifier on graphs with scarce and noisy node labels. We proposed a novel RTGNN framework which performs explicit noise governance with supplemental supervision. More specifically, we carefully classified labeled nodes into clean and noisy ones and adopted reinforcement supervision to correct noisy labels. We also created pseudo labels to provide extra training signals. Moreover, we leveraged consistency regularization to prevent overfitting to noise. By doing so, RTGNN could enable sufficient learning from clean labels while limiting the impact of noisy ones. Comprehensive experiments on three widely-used datasets were conducted, and the results showed superior performance of our RTGNN under two types of noise and varying noise rates. We expect that our work would inspire more studies on robust semi-supervised learning on graphs.

ACKNOWLEDGMENTS

This research was partially supported by National Key R&D Program of China under grant No. 2018AAA0102102, National Natural Science Foundation of China under grants No. 62132017 and No. 62106218, D. Z. Chen’s research was supported in part by NSF Grant CCF-1617735.

REFERENCES

- [1] Firoj Alam, Shafiq Joty, and Muhammad Imran. 2018. Graph based semi-supervised learning with convolution neural networks to classify crisis related tweets. In *12th International AAAI Conference on Web and Social Media*.
- [2] Devansh Arpit, Stanislaw Jastrzębski, Nicolas Ballas, David Krueger, Emmanuel Bengio, Maxinder S Kanwal, Tegan Maharaj, Asja Fischer, Aaron Courville, Yoshua Bengio, et al. 2017. A closer look at memorization in deep networks. In *International Conference on Machine Learning*. PMLR, 233–242.
- [3] Joan Bruna, Wojciech Zaremba, Arthur Szlam, and Yann LeCun. 2014. Spectral Networks and Locally Connected Networks on Graphs. In *2nd International Conference on Learning Representations, ICLR*.
- [4] Yu Chen, Lingfei Wu, and Mohammed J Zaki. 2019. Deep iterative and adaptive learning for graph neural networks. *arXiv preprint arXiv:1912.07832* (2019).
- [5] Luca Cosmo, Anees Kazi, Seyed-Ahmad Ahmadi, Nassir Navab, and Michael Bronstein. 2020. Latent-graph learning for disease prediction. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 643–653.
- [6] Enyan Dai, Charu Aggarwal, and Suhang Wang. 2021. NRGNN: Learning a Label Noise Resistant Graph Neural Network on Sparsely and Noisily Labeled Graphs. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 227–236.
- [7] Enyan Dai, Wei Jin, Hui Liu, and Suhang Wang. 2022. Towards Robust Graph Neural Networks for Noisy Graphs with Sparse Labels. In *WSDM'22: The 15th ACM International Conference on Web Search and Data Mining*.
- [8] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. 2016. Convolutional neural networks on graphs with fast localized spectral filtering. *Advances in Neural Information Processing Systems* 29 (2016).
- [9] Bahare Fatemi, Layla El Asri, and Seyed Mehran Kazemi. 2021. SLAPS: Self-Supervision Improves Structure Learning for Graph Neural Networks. *Advances in Neural Information Processing Systems* 34 (2021).
- [10] Aritra Ghosh, Himanshu Kumar, and PS Sastry. 2017. Robust loss functions under label noise for deep neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 31.
- [11] Jacob Goldberger and Ehud Ben-Reuven. 2017. Training deep neural-networks using a noise adaptation layer. In *5th International Conference on Learning Representations (ICLR) 2017*.
- [12] Jonathan Halcrow, Alexandru Mosoi, Sam Ruth, and Bryan Perozzi. 2020. GraLe: Designing networks for graph learning. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2523–2532.
- [13] Will Hamilton, Zhitaoying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. *Advances in Neural Information Processing Systems* 30 (2017).
- [14] Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor Tsang, and Masashi Sugiyama. 2018. Co-teaching: Robust training of deep neural networks with extremely noisy labels. *Advances in Neural Information Processing Systems* 31 (2018).
- [15] Song Han, Huizi Mao, and William J. Dally. 2016. Deep Compression: Compressing Deep Neural Network with Pruning, Trained Quantization and Huffman Coding. In *4th International Conference on Learning Representations (ICLR)*, Yoshua Bengio and Yann LeCun (Eds.).
- [16] Jinchuan Huang, Lie Qu, Rongfei Jia, and Binqiang Zhao. 2019. O2U-net: A simple noisy label detection approach for deep neural networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 3326–3334.
- [17] Lu Jiang, Zhengyuan Zhou, Thomas Leung, Li-Jia Li, and Li Fei-Fei. 2018. MentorNet: Learning data-driven curriculum for very deep neural networks on corrupted labels. In *International Conference on Machine Learning*. PMLR, 2304–2313.
- [18] Wei Jin, Yao Ma, Xiaorui Liu, Xianfeng Tang, Suhang Wang, and Jiliang Tang. 2020. Graph structure learning for robust graph neural networks. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 66–74.
- [19] Anees Kazi, Luca Cosmo, Seyed-Ahmad Ahmadi, Nassir Navab, and Michael Bronstein. 2022. Differentiable graph module (DGM) for graph convolutional networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2022).
- [20] Thomas N Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907* (2016).
- [21] Thomas N Kipf and Max Welling. 2016. Variational graph auto-encoders. *arXiv preprint arXiv:1611.07308* (2016).
- [22] Christos Louizos, Max Welling, and Diederik P. Kingma. 2018. Learning Sparse Neural Networks through L_0 Regularization. In *6th International Conference on Learning Representations (ICLR)*.
- [23] Xingjun Ma, Hanxun Huang, Yisen Wang, Simone Romano, Sarah Erfani, and James Bailey. 2020. Normalized loss functions for deep learning with noisy labels. In *International Conference on Machine Learning*. PMLR, 6543–6553.
- [24] Eran Malach and Shai Shalev-Shwartz. 2017. Decoupling “when to update” from “how to update”. *Advances in Neural Information Processing Systems* 30 (2017).
- [25] Miller McPherson, Lynn Smith-Lovin, and James M Cook. 2001. Birds of a Feather: Homophily in Social Networks. *Annual Review of Sociology* 27, 1 (2001), 415–444.
- [26] Duc Tam Nguyen, Chaithanya Kumar Mummadi, Thi-Phuong-Nhung Ngo, Thi Hoai Phuong Nguyen, Laura Beggel, and Thomas Brox. 2020. SELF: Learning to Filter Noisy Labels with Self-Ensembling. In *8th International Conference on Learning Representations (ICLR)*.
- [27] Hoang NT, Choong Jun Jin, and Tsuyoshi Murata. 2019. Learning graph neural networks with noisy labels. *arXiv preprint arXiv:1905.01591* (2019).
- [28] Shirui Pan, Ruiqi Hu, Guodong Long, Jing Jiang, Lina Yao, and Chengqi Zhang. 2018. Adversarially regularized graph autoencoder for graph embedding. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*. 2609–2615.
- [29] 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 1944–1952.
- [30] Geoff Pleiss, Tianyi Zhang, Ethan Elenberg, and Kilian Q Weinberger. 2020. Identifying mislabeled data using the area under the margin ranking. *Advances in Neural Information Processing Systems* 33 (2020), 17044–17056.
- [31] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. 2017. Graph attention networks. *arXiv preprint arXiv:1710.10903* (2017).
- [32] Daixin Wang, Jianbin Lin, Peng Cui, Quanhuai Jia, Zhen Wang, Yanming Fang, Quan Yu, Jun Zhou, Shuang Yang, and Yuan Qi. 2019. A Semi-supervised Graph Attentive Network for Financial Fraud Detection. In *2019 IEEE International Conference on Data Mining (ICDM)*. IEEE, 598–607.
- [33] Deng-Bao Wang, Yong Wen, Lujia Pan, and Min-Ling Zhang. 2021. Learning from noisy labels with complementary loss functions. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 10111–10119.
- [34] Yisen Wang, Xingjun Ma, Zaiyi Chen, Yuan Luo, Jinfeng Yi, and James Bailey. 2019. Symmetric cross entropy for robust learning with noisy labels. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 322–330.
- [35] Hongxin Wei, Lei Feng, Xiangyu Chen, and Bo An. 2020. Combating noisy labels by agreement: A joint training method with co-regularization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 13726–13735.
- [36] Jun Wu, Jingrui He, and Jiejun Xu. 2019. DEMO-Net: Degree-specific graph neural networks for node and graph classification. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 406–415.
- [37] Xiaobo Xia, Tongliang Liu, Bo Han, Mingming Gong, Jun Yu, Gang Niu, and Masashi Sugiyama. 2021. Sample selection with uncertainty of losses for learning with noisy labels. *arXiv preprint arXiv:2106.00445* (2021).
- [38] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. 2018. How powerful are graph neural networks? *arXiv preprint arXiv:1810.00826* (2018).
- [39] Zhilin Yang, William Cohen, and Ruslan Salakhudinov. 2016. Revisiting semi-supervised learning with graph embeddings. In *International Conference on Machine Learning*. PMLR, 40–48.
- [40] Qing Yu, Atsushi Hashimoto, and Yoshitaka Ushiku. 2021. Divergence Optimization for Noisy Universal Domain Adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2515–2524.
- [41] Xingrui Yu, Bo Han, Jiangchao Yao, Gang Niu, Ivor Tsang, and Masashi Sugiyama. 2019. How does disagreement help generalization against label corruption?. In *International Conference on Machine Learning*. PMLR, 7164–7173.
- [42] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. 2017. Understanding deep learning requires rethinking generalization. In *5th International Conference on Learning Representations (ICLR)*.
- [43] Ying Zhang, Tao Xiang, Timothy M Hospedales, and Huchuan Lu. 2018. Deep mutual learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 4320–4328.
- [44] Zhilu Zhang and Mert Sabuncu. 2018. Generalized cross entropy loss for training deep neural networks with noisy labels. *Advances in Neural Information Processing Systems* 31 (2018).
- [45] Tong Zhao, Yozen Liu, Leonardo Neves, Oliver Woodford, Meng Jiang, and Neil Shah. 2021. Data augmentation for graph neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 11015–11023.
- [46] Haoyi Zhou, Shanghang Zhang, Jieqi Peng, Shuai Zhang, Jianxin Li, Hui Xiong, and Wancai Zhang. 2021. Informer: Beyond efficient transformer for long sequence time-series forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 11106–11115.