

Distance-wise Prototypical Graph Neural Network for Imbalanced Node Classification

Yu Wang
yu.wang.1@vanderbilt.edu
Vanderbilt University

Charu Aggarwal
charu@us.ibm.com
IBM T.J. Watson Research Center

Tyler Derr
tyler.derr@vanderbilt.edu
Vanderbilt University

ABSTRACT

Recent years have witnessed the significant success of applying graph neural networks (GNNs) in learning effective node representations for classification. However, current GNNs are mostly built under the balanced data-splitting, which is inconsistent with real-world networks where the number of training nodes can be extremely imbalanced among the classes. Thus, directly utilizing current GNNs on imbalanced data would generate coarse representations of nodes in minority classes. This therefore portends the importance of developing effective GNNs for handling imbalanced graph data. In this work, we propose a novel Distance-wise Prototypical Graph Neural Network (DPGNN), which utilizes a class prototype-driven training to balance the training loss between majority and minority classes and then leverages distance metric learning to differentiate the contributions of different dimensions of representations and fully encode the relative position of each node to each class prototype. Moreover, we design a new imbalanced label propagation mechanism to derive extra supervision from unlabeled nodes and employ self-supervised learning to smooth representations of adjacent nodes while separating inter-class prototypes. Comprehensive node classification experiments and parameter analysis on multiple networks are conducted and the proposed DPGNN almost always significantly outperforms all other baselines. The implementation of DPGNN is available at <https://github.com/YuWVandy/DPGNN>.

KEYWORDS

Imbalanced node classification, graph neural networks

ACM Reference Format:

Yu Wang, Charu Aggarwal, and Tyler Derr. 2022. Distance-wise Prototypical Graph Neural Network for Imbalanced Node Classification. In *International Workshop on Mining and Learning with Graphs*. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Graph neural networks (GNNs) have become one of the most promising paradigms in graph representation learning for node classification [14], where a classifier is trained by labeled nodes and then used for categorizing labels of all remaining nodes. Although many GNN variants have been proposed to complete this

task [10, 18], prevailing prior works follow a (semi-)supervised setting where labeled nodes are assumed to be balanced among different classes [35]. However, this requires sufficient balanced labeled data, which is over-idealized and inconsistent with reality.

In many real-world networks, the class distribution of labeled nodes is inherently skewed [34] where a large portion of classes (minority classes) only contain a limited number of labeled nodes (minority nodes) while few classes (majority classes) contain enough labeled nodes (majority nodes). Since most GNNs are designed without considering the potential of class imbalance, directly using them on imbalanced dataset would undermine the learned representations of minority nodes. The imbalanced node classification with GNNs naturally inherits existing challenges of deep learning in imbalanced classification: the inclination to learning towards majority classes [13] and the catastrophic forgetting of previous learned instances in minority classes [29]. First, deep learning models improve representations for completing target tasks via backpropagation from the training loss. However, in the class imbalanced scenario, the main component of the training loss comes from majority classes and thus the gradient is dominated by majority classes such that the model is updated towards behaving significantly better on majority classes than minority ones. Second, deep learning models are notorious for demanding big data for updating parameters [29], which limits their ability to learn from minority classes that may only have a few training instances in the imbalanced setting. In addition, over-smoothing [17] that occurs as a general issue in GNNs would become even worse in the imbalanced setting: representations of minority nodes would become similar to the majority ones and deviate from their own spectrum due to the imbalance bias introduced in the message passing.

Traditional methods for handling imbalance are either augmenting data via under(over)-sampling [4], or assigning weights to adjust the portion of training loss from different classes [27]. Most of these methods are proposed for non-graph structured data. DR-GCN [25] is the pioneer to explore node imbalance classification by adversarial training each class distribution and enforce the consistency between the labeled and unlabeled data distributions. However, the adversarial training may yield unrepresentative minority class distribution given few minority nodes [8]. Further, RECT [30] is proposed by leveraging topological regularization to derive extra supervision for minority nodes while it is designed for completely-imbalanced data and its performance on partially imbalanced data is unclear. GraphSMOTE [34] generalizes SMOTE [3] to the graph domain by pre-training an edge generator and hence adding relational information for the new synthetic nodes from SMOTE. However, the computation of calculating the similarity between all pairs of nodes and pre-training the edge generator is extremely heavy.

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.

MLG '22, August 15, 2022, Washington, D.C.

© 2022 Association for Computing Machinery.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM...\$15.00

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

To tackle the aforementioned challenges of imbalanced node classification, we present Distance-wise Prototypical Graph Neural Network (DPGNN), which first applies class prototype-driven training to balance the training loss of different classes and then leverages distance metric learning to differentiate the contribution of each dimension of distance from each query node to all class prototypes. Ultimately the classification is performed by comparing the similarity of the learned distance metric representations of the nodes with the ones of class prototypes. Additionally, DPGNN introduces a novel imbalance label propagation scheme to augment training data and employs self-supervised learning to smooth representations of adjacent nodes while separating inter-class prototypes. The main contributions are summarized as follows:

- We construct a balanced training scheme inspired by episodic training and further introduce distance metric learning to better capture node distances to class prototypes.
- We design a new imbalanced label propagation scheme to augment the training data and employ self-supervised learning to further improve the distance metric representations.
- We perform extensive imbalanced node classification experiments on real-world datasets across various levels of class imbalance with detailed parameter analysis to corroborate the effectiveness of our model.

In Section 2, we introduce related work. Section 3 introduces necessary notations and the imbalanced node classification problem. In Section 4, we present our proposed DPGNN model. We evaluate our framework by performing extensive experiments and provide clear interpretations in Section 5. Finally, we conclude and discuss future work in Section 6.

2 RELATED WORK

2.1 Class Imbalance Problem

Class imbalance exists in many real-world applications [12], where classes with more (or less) training instances are termed as majority (or minority) classes. The imbalance in the number of training instances among different classes significantly affects the performance of supervised learning and hence has become a classical research direction [13]. Generally, the approaches against this problem are summarized into three levels, i.e., data-level, algorithm-level and hybrid [16]. Data-level methods aim to improve the data by balancing the training instances (e.g., up(down)-sampling [4]), whereas algorithm-level methods [27] attempt to improve the training process by modifying the learning algorithm such as re-weighting. In this work, we research on class imbalance in the graph domain. In the algorithm-level, our work balances the training loss by comparing labeled nodes with each class prototypes. In the data-level, we propose an imbalanced label propagation to obtain additional minority samples. Thus, our DPGNN is a hybrid approach.

2.2 Graph Neural Networks

GNNs have achieved unprecedented success, but most prior work fails to consider the class imbalance problem, which unfortunately widely exists in real-world applications [30, 34]. RECT [30] was developed merging a GNN and proximity-based embedding component for the completely-imbalanced setting (i.e., where some classes can even have no labeled nodes). DR-GCN [25] explored node imbalance classification by adversarial training each class distribution

and enforce the consistency between the labeled and unlabeled data distributions. Unfortunately, the adversarial trained generator and discriminator maybe overfitted to few minority nodes and hence loses generalizability [8]. More recently, GraphSMOTE [34] was designed using a GNN encoder to learn node embeddings and an extra edge generator to generate edges connecting synthetic minority nodes. However, the model is time-consuming and somewhat learned in an ad-hoc fashion. Motivated by this and to provide a grounded approach for imbalanced node classification, we design a novel GNN-based framework that utilizes class prototypes to balance the training loss and distance metric learning to fully encode the relative position of each node to each class prototype.

3 PROBLEM STATEMENT

We denote an attributed graph by $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{X})$ where $\mathcal{V} = \{v_1, \dots, v_n\}$ is the set of n nodes, \mathcal{E} is the set of m edges with e_{ij} being the edge between nodes v_i and v_j , and $\mathbf{X} = [\mathbf{x}_1^\top, \dots, \mathbf{x}_n^\top] \in \mathbb{R}^{n \times d}$ is the node feature matrix with \mathbf{x}_i indicating the features of node v_i . The network topology is described by its adjacency matrix $\mathbf{A} \in \{0, 1\}^{n \times n}$, where $\mathbf{A}_{ij} = 1$ denotes an edge between nodes v_i and v_j , and $\mathbf{A}_{ij} = 0$ otherwise. The diagonal matrix of node degrees is denoted as $\mathbf{D} \in \mathbb{R}^{n \times n}$, where $\mathbf{D}_{ii} = \sum_j \mathbf{A}_{ij}$ calculates the degree of the node v_i . We let $\tilde{\mathbf{A}} = \mathbf{A} + \mathbf{I}$ represent the adjacency matrix with added self-loops and similarly let $\tilde{\mathbf{D}} = \mathbf{D} + \mathbf{I}$. Then the normalized adjacency matrix can be defined as $\tilde{\mathbf{A}} = \tilde{\mathbf{D}}^{-0.5} \mathbf{A} \tilde{\mathbf{D}}^{-0.5}$. The neighborhood node set of the center node v_i is given by $\mathcal{N}_i = \{v_j | e_{ij} \in \mathcal{E}\}$. Finally, \mathcal{C} is the set of C classes and $\mathbf{Y} \in \mathbb{R}^{n \times C}$ is the one-hot node label matrix. Given these notations, the imbalanced node classification can be mathematically formalized as:

Given an attributed network $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{X})$ with labels for a subset of nodes $\mathcal{V}_l \subset \mathcal{V}$ that are imbalanced among the C classes, we aim to learn a node classifier $f : f(\mathcal{V}, \mathcal{E}, \mathbf{X}) \rightarrow \mathbf{Y}$ that can work well for both majority and minority classes.

4 THE PROPOSED FRAMEWORK

In this section, we present our proposed Distance-wise Prototypical Graph Neural Network (DPGNN) that is able to solve the challenges imposed on deep graph learning when given imbalanced training data. The framework is shown in Figure 1, including four main components: (a) class prototype-driven balanced training, (b) distance metric learning, (c) imbalanced label propagation, and (d) self-supervised learning. Specifically, we borrow the idea of episodic training from prototypical networks [26] to first balance the training loss across different classes via prototype-driven balanced training (Figure 1(a)). However, unlike traditional prototypical networks where labels of query nodes are assigned based on their nearest class prototypes, we further employ distance metric learning to construct another distance metric space that can differentiate different distance dimensions and fully characterize the position of each query node relative to all class prototypes (Figure 1(b)). In order to fully incorporate the graph topology information in alleviating the imbalance problem, we also design a novel imbalanced label propagation scheme and two self-supervised components to aid in learning high-quality distance metric representations (Figure 1(c) and 1(d), respectively). Next we describe each component in detail.

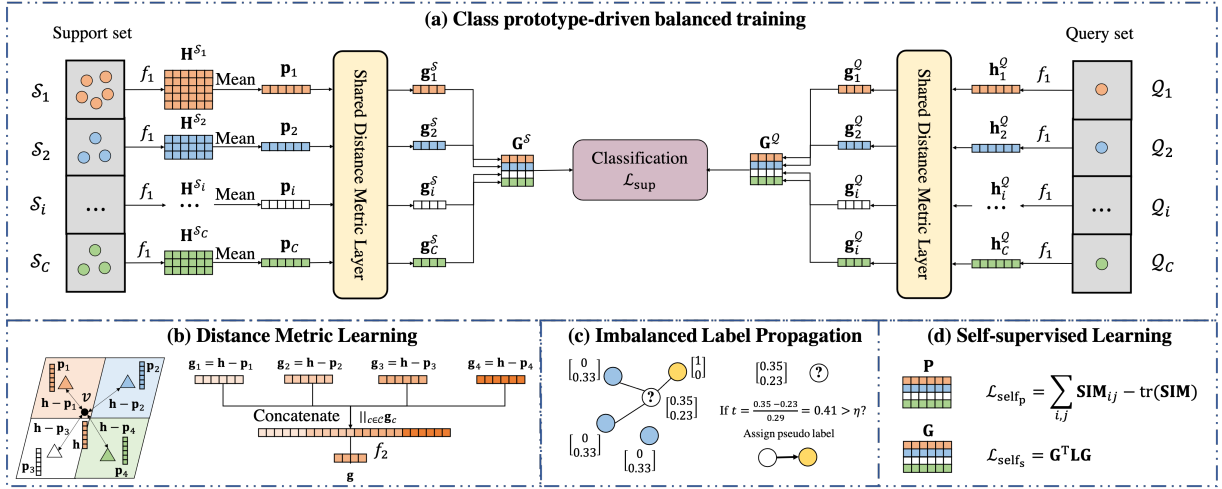


Figure 1: Overview of the Distance-wise Prototypical Graph Neural Network, with four main components: (a) class prototype-driven balanced training, (b) distance metric learning, (c) imbalanced label propagation, and (d) self-supervised learning.

4.1 Class Prototype-driven Balanced Training

In order to balance the training loss from majority and minority classes, we leverage the idea of episodic training by sampling support and query sets, and then calculating representations of prototypes for each class. In order to avoid using the original sparse and high-dimensional node features and to allow learned complex feature transformation, we first apply a GNN-based encoder $f_l: \mathbb{R}^{n \times d} \times \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^{n \times d'}$ to obtain d' -dimensional node representations $\mathbf{H} \in \mathbb{R}^{n \times d'}$. Most GNN-based encoders f_l can be decomposed into two components: neighborhood propagation and feature transformation, which can be generally formalized as:

$$\mathbf{h}_i^l = \text{TRAN}^l(\text{AGGR}^l(\mathbf{h}_i^{l-1}, \text{PROP}^l(\{\mathbf{h}_j^{l-1} | v_j \in \mathcal{N}_i\}))), \quad (1)$$

where in each layer l , the representations $\{\mathbf{h}_j^{l-1} | v_j \in \mathcal{N}_i\}$ of the neighbors are first propagated via PROP^l to node v_i and aggregated with its own representation \mathbf{h}_i^{l-1} by AGGR^l , then the combined representation is further transformed by TRAN^l to output the representation \mathbf{h}_i^l of node v_i after layer l of the GNN.

With the learned node embeddings \mathbf{H} coming from the GNN-based encoder's last layer, we aim to compute representations of class prototypes $\mathbf{P} \in \mathbb{R}^{C \times d'}$. During each training epoch, the model is fed with an episode sampled from the labeled training nodes \mathcal{V}_l , which is further divided into support sets $\mathcal{S} = \{\mathcal{S}_c | c \in \mathcal{C}\}$ and query sets $\mathcal{Q} = \{\mathcal{Q}_c | c \in \mathcal{C}\}$ of each class. The query sets only contain one training sample from each class and provide a balanced way to train the encoder. The support sets are formed by grouping the remaining training samples from each class except the query samples from the query sets, which serves as anchors to characterize the class prototypes of each class. Based on Prototypical Networks [26], we define class prototypes to be closely surrounded by nodes of the same class, such that they can precisely represent their class. More specifically, the prototype \mathbf{p}_c of class c is computed by:

$$\mathbf{p}_c = \text{PROTO}(\mathbf{H}^{\mathcal{S}_c}) = \text{PROTO}(\{\mathbf{h}_i | v_i \in \mathcal{S}_c\}), c \in \mathcal{C}, \quad (2)$$

where PROTO is the prototype computation which calculates the representation of a class prototype \mathbf{p}_c based on node representations

$\mathbf{H}^{\mathcal{S}_c}$ that come from the support set \mathcal{S}_c . For instance, in the vanilla Prototypical Network [26], the mean-pooling is employed here:

$$\mathbf{p}_c = \frac{1}{|\mathcal{S}_c|} \sum_{v_i \in \mathcal{S}_c} \mathbf{h}_i, c \in \mathcal{C}. \quad (3)$$

Applying mean-pooling to calculate class prototypes assumes that each class can be represented using only one class prototype, which might be underrepresentative when a unimodal distribution assumption is violated. Thus, multi-prototypes to represent each class could be used, such as replacing mean-pooling with another permutation-invariant function (e.g., K-means clustering [11]), or even iteratively learn multi-prototypes according to the complexity of the class distribution [2]. We leave this as one future direction.

Prototypes serve as representatives of each class and can be used to classify query nodes by selecting their nearest prototypes. However, directly employing a softmax over distances to prototypes in the embedding space to obtain the class probability distribution [7, 26, 29] will make different dimensions of distance contribute equally to the classification and in high dimensional data pairwise Euclidean distances tend to converge [1]. Furthermore, using only embedding distance of a node to the nearest prototype omit its embedding distance to all other non-nearest prototypes that may encode extra information about the node's class. Thus, instead of classifying query nodes directly based on their nearest prototype [7, 26], we devise a distance metric layer to project nodes from the original embedding space to another distance metric space, where query nodes are classified by comparing their learned distance metric representations with the ones of class prototypes.

4.2 Distance Metric Learning

In order to project each node from the original embedding space to the distance metric space, we first concatenate the difference of its embedding from each class prototype. Next we apply a linear transformation on top of that, to pay different attention on each original distance dimension and adaptively extract useful distance information. Figure 1(b) demonstrates the detailed procedure of computing distance metric representation \mathbf{g} of node v .

Given a node v with embedding $\mathbf{h} \in \mathbb{R}^{d'}$, we calculate its distance metric representation $\mathbf{g}_c \in \mathbb{R}^{d'}$ to each class prototype c as:

$$\mathbf{g}_c = \mathbf{h} - \mathbf{p}_c, c \in \mathcal{C}, \quad (4)$$

where $\mathbf{h} - \mathbf{p}_c$ calculates the difference of the embedding between each node and each class prototype. For each node, only considering its embedding difference to one class prototype cannot fully locate its position. Therefore, we concatenate the difference of the node's embedding to all class prototypes $\|\mathbf{g}_c\|_{c \in \{1, \dots, C\}}$ and further apply a linear transformation $f_2: \mathbb{R}^{d'C} \rightarrow \mathbb{R}^{d''}$ to pay different levels of attention to different dimensions of the embedding difference and adaptively extract useful embedding difference information.

$$\mathbf{g} = f_2(\|\mathbf{g}_c\|_{c \in \mathcal{C}}). \quad (5)$$

The distance metric representation \mathbf{g} encodes the distance information of the node v to all class prototypes, which as a result precisely captures its relative position to all class prototypes. In order to use these distance metric representations to prototypes for reference to classify query nodes, we feed the representations of prototypes and query nodes $\mathbf{p}_c, \mathbf{h}_c^Q$, for $c \in \mathcal{C}$ into the shared distance metric layer to learn their distance metric representations $\mathbf{g}_c^S, \mathbf{g}_c^Q$, for $c \in \mathcal{C}$. Then we stack them across all classes to compute the distance metric representations of prototypes and query nodes:

$$\mathbf{G}^{\mathcal{M}} = \text{stack}(\mathbf{g}_c^{\mathcal{M}} | c \in \mathcal{C}), \mathcal{M} \in \{\mathcal{Q}, \mathcal{S}\} \quad (6)$$

Next, the predicted class distribution for each query node is:

$$\mathbf{F} = \text{softmax}(\mathbf{G}^{\mathcal{Q}}(\mathbf{G}^{\mathcal{S}})^{\top}), \quad (7)$$

where \mathbf{F}_c gives the predicted probability distribution of class c 's query node over all classes, which is then used to calculate the supervised classification loss:

$$\mathcal{L}_{\text{class}} = \frac{1}{C} \sum_{c=1}^C \ell(\mathbf{F}_c, c), \quad (8)$$

where $\ell(\cdot, \cdot)$ is a loss function to measure the difference between predictions and ground-truth labels, such as cross entropy.

4.3 Imbalanced Label Propagation

Network homophily [21] assumes that connected nodes tend to share similar attributes or belong to the same class, which is commonly the case in various real-world networks [36]. Thus, it can be naturally harnessed to augment the training data especially for increasing the minority training nodes. Motivated by this, we perform imbalanced label propagation to annotate node labels based on their neighboring nodes' labels and filter out unconfident ones by computing their topology information gain [5].

Since we only have access to the class information of labeled nodes \mathcal{V}_l and to balance the voting effect between the majority and the minority classes, for unlabeled nodes we mask their one-hot labels by filling zero vector $\mathbf{0} \in \mathbb{R}^C$, and for labeled nodes we multiply their one-hot labels by the weighting factor γ_i :

$$\tilde{\mathbf{Y}}_i = \begin{cases} \gamma_i \mathbf{Y}_i, & v_i \in \mathcal{V}_l \\ \mathbf{0}, & v_i \in \mathcal{V} \setminus \mathcal{V}_l, \end{cases} \text{ where } \gamma_i = \frac{|\mathcal{V}_l|}{\sum_{v_j \in \mathcal{V}_l} \mathbf{Y}_{j\phi(v_i)}} \quad (9)$$

computes the inverse ratio of the number of labeled nodes in the node v_i 's class $\phi(v_i)$, $\sum_{v_j \in \mathcal{V}_l} \mathbf{Y}_{j\phi(v_i)}$, to the total number of labeled nodes $|\mathcal{V}_l|$. Then we propagate this unbiased label distribution $\tilde{\mathbf{Y}}$ to their neighbors that are at most k -hops away by considering

k^{th} -order adjacency matrix as follows:

$$\hat{\mathbf{Y}} = \hat{\mathbf{A}}^k \tilde{\mathbf{Y}}. \quad (10)$$

Note that $\hat{\mathbf{A}}^k \tilde{\mathbf{Y}}$ can be computed efficiently by applying power iteration sequentially from right to left [15].

However, this leads to another issue that $\hat{\mathbf{Y}}$ cannot be used directly in class prototype-driven balance training to sample support and query sets since $\hat{\mathbf{Y}}$ is a soft-label (instead of having hard label assignments). Moreover, nodes that fall in the topological boundary between different classes may possess noisy label distribution and mislead the training process. Therefore, we utilize the idea of topological information gain (TIG) [5] to filter out nodes with weak and obscure label distributions. TIG describes the task information effectiveness that the node obtains from labeled source along the network topology [5]. We regard the maximum entry of the soft-label $\hat{\mathbf{Y}}_i$ to be the most possible class type that the node v_i can be and the other entries as confusing information. Then the topological information gain \mathbf{t}_i for node v_i is calculated as:

$$\mathbf{t}_i = \frac{\max(\hat{\mathbf{Y}}_i) - \frac{(\sum_{c \in \mathcal{C}} \hat{\mathbf{Y}}_{ic} - \max(\hat{\mathbf{Y}}_i))}{C-1}}{\frac{1}{C} \sum_{c \in \mathcal{C}} \hat{\mathbf{Y}}_{ic}}. \quad (11)$$

A high \mathbf{t}_i means the label distribution of node v_i is sharp/strong and thus it lies in the cluster of nodes in class $\text{argmax}(\hat{\mathbf{Y}}_i)$. By network homophily, its label is also conjectured to be $\text{argmax}(\hat{\mathbf{Y}}_i)$. As such, we generate the hard pseudo label of non-training node v_i by binary thresholding its topological information gain \mathbf{t}_i as:

$$\check{\mathbf{Y}}_{ic} = \begin{cases} 1, & \text{if } \mathbf{t}_i > \eta \text{ and } c = \text{argmax}(\hat{\mathbf{Y}}_i) \\ 0, & \text{if } \mathbf{t}_i \leq \eta \text{ or } c \neq \text{argmax}(\hat{\mathbf{Y}}_i) \end{cases}, v_i \in \mathcal{V} \setminus \mathcal{V}_l, \quad (12)$$

where η is a hyperparameter that controls the trade-off between the quality and the number of the augmented labels. Higher η leads to precise labels with sharp distribution and thus guarantee the labeling quality but less nodes could be augmented. Lower η leads to more imprecise and noisy labels with even distribution. Note that for the initial labeled nodes $v_i \in \mathcal{V}_l$, we still utilize their original labels and thus the final labels we use for sampling support and query sets in class prototype-driven balanced training:

$$\tilde{\mathbf{Y}}_i = \begin{cases} \mathbf{Y}_i & v_i \in \mathcal{V}_l \\ \check{\mathbf{Y}}_i & v_i \in \mathcal{V} \setminus \mathcal{V}_l \end{cases} \quad (13)$$

Details of hyperparameter tuning η are provided in Section 5.4.

4.4 Self-Supervised Learning (SSL)

Although node embeddings computed from the GNN-based encoder f_1 embed network topology information, the prototype computation PROTO in Eq. (2) and the distance metric learning in Eq. (4) do not consider network topology. Inspired by the intuition that different class prototypes should have different representations and adjacent nodes should have similar distance metric representations, we design two GNN self-supervised learning (SSL) [20] pretext tasks to emphasize the topological information learned by DPGNN.

To ensure that different prototypes have different representations, we minimize the representation similarity of prototypes from different classes as follows:

$$\mathcal{L}_{\text{sslp}} = \sum_{i=1}^C \sum_{j=1}^C (\text{SIM}_{ij}) - \text{tr}(\text{SIM}), \text{ where } \text{SIM}_{ij} = \frac{\mathbf{p}_i^{\top} \mathbf{p}_j}{\|\mathbf{p}_i\| \|\mathbf{p}_j\|} \quad (14)$$

To smooth the distance metric representations between adjacent nodes and reinforce the graph structure in the learned representations, we adopt the following objective function:

$$\mathcal{L}_{\text{ssl}_s} = \mathbf{G}^T \mathbf{L} \mathbf{G}, \quad (15)$$

where \mathbf{G} is the distance metric representations by transforming the concatenated embedding distance between each node to each class prototype (i.e., Eq. (4) and Eq. (5)), and \mathbf{L} is the Laplacian matrix of the underlying network. If we adopt the normalized Laplacian matrix, i.e., $\mathbf{L} = \mathbf{I} - \hat{\mathbf{A}}$, Eq. (15) can be rewritten as:

$$\mathcal{L}_{\text{ssl}_s} = \sum_{v_i \in \mathcal{V}} \sum_{v_j \in \mathcal{N}_i} \left(\frac{\mathbf{g}_i}{\sqrt{d_i}} - \frac{\mathbf{g}_j}{\sqrt{d_j}} \right)^2. \quad (16)$$

It is clear that $\mathcal{L}_{\text{ssl}_s}$ is small when adjacent nodes share similar distance metric representations to class prototypes.

Collecting one supervised classification loss and two SSL loss, the overall objective function is stated as:

$$\min \mathcal{L} = \mathcal{L}_{\text{class}} + \lambda_1 \mathcal{L}_{\text{ssl}_p} + \lambda_2 \mathcal{L}_{\text{ssl}_s}, \quad (17)$$

where λ_1 and λ_2 are two hyperparameters that control the contribution of the two SSL losses (i.e., $\mathcal{L}_{\text{ssl}_p}$ and $\mathcal{L}_{\text{ssl}_s}$) in addition to the supervised classification loss $\mathcal{L}_{\text{class}}$.

4.5 Complexity Analysis

In comparison to vanilla GNN-based encoders, additional computational requirements come from three components: distance metric learning, imbalanced label propagation, and self-supervised learning. For distance metric learning, the calculation of the pairwise embedding difference between each node and each class prototype requires $O(|\mathcal{V}|C)$ time complexity if implemented naively. Typically the class number C is multiple orders of magnitude less than the node number $|\mathcal{V}|$ in a network and therefore can be treated as a constant, which leads to linear time complexity $O(|\mathcal{V}|)$. For uncommon networks that have an extremely large number of classes with the same magnitude as the size of the network, we could select sub-classes C_{sub} as anchors and approximate the distance metric representations by considering distance to these anchors rather than all classes [6, 19, 31]. In imbalanced label propagation, the most computational part comes from propagating labels in Eq. (10), which can be completed efficiently by applying power iteration from the edge view in $O(k|\mathcal{E}|)$ compared to $O(|\mathcal{V}|^3)$ for matrix calculation of $\hat{\mathbf{A}}^k$. Among the two SSL components, the heaviest computation comes from smoothing distance metric embeddings of adjacent nodes in Eq. (15), which can be calculated from the edge view as Eq. (16); thus, is linear with the number of edges $O(|\mathcal{E}|)$.

For the model complexity, apart from the parameters of GNN-based encoder, additional parameters of DPGNN come from the linear transformation f_2 , which are \mathbf{W}^{f_2} and \mathbf{b}^{f_2} . Hence, compared with vanilla GNN-based encoders, the overall additional parameters are $O(d'C \times d'')$ where $d'C$ is the dimension of concatenated distance metric representations to all class prototypes $\|_{c \in \{1, \dots, C\}} \mathbf{g}_c$ and d'' is the dimension of distance metric representation \mathbf{g} (as in Eq. (5)). Typically, if the number of classes C is small or if we only select some sub-classes as anchors, $d'C$ will be far less than the original high-dimensional node attributes d . Therefore, the extra model complexity $O(d'C \times d'')$ could practically be ignored compared to the complexity of the GNN-based encoder, which has $O(d \times d')$.

Table 1: Basic dataset statistics.

Networks	Nodes	Edges	Features	Classes	Homophily	Type
Cora	2,708	5,429	1,433	7	0.81	Citation
Citeseer	3,327	4,732	3,703	6	0.74	Citation
Pubmed	19,717	44,338	500	3	0.80	Citation
Cora-ML	2,995	8,416	2,879	7	0.79	Citation
DBLP	17,716	105,734	1,639	4	0.83	Citation
Amazon Computers	13,381	245,778	767	10	0.78	Online Product
Amazon Photo	7,487	119,043	745	8	0.83	Online Product
Twitch PT	1,912	64,510	128	2	0.58	Social

5 EXPERIMENTAL RESULTS

In this section, we conduct extensive experiments on imbalanced node classification to evaluate the effectiveness of DPGNN. In particular, we target to answer the following three questions:

- **Q1:** How effective is DPGNN compared to other baselines on imbalanced node classification under different imbalance ratio?
- **Q2:** How do different components of DPGNN contributes to performance improvement in imbalanced node classification?
- **Q3:** How does the threshold η in imbalanced label propagation affect the performance of DPGNN?

5.1 Experiment Settings

5.1.1 Datasets. We experiment on widely-adopted citation networks [32], Amazon product networks [24], and an online social network [23]. Table 1 presents the basic network statistics for these datasets. In the five citation networks and the online social network, class distributions are relatively balanced, so we use an imitative imbalanced setting: we choose the first 5, 4, 2, 5, 3, 1 class(es) as minority and down-sample their training nodes to 2 compared to 20 for other majority class(es), which creates an imbalanced class distribution with imbalance ratio 10. For the Amazon product networks whose class distributions are genuinely imbalanced, we use their original class ratios and set the total training nodes as 50 and 30, respectively. For validation and testing sets, 500 and 1000 nodes are selected respectively for all 8 datasets, which is commonly employed in the literature [14, 28]. This setting is used throughout the paper unless otherwise stated.


5.1.2 Baselines. To evaluate the effectiveness of the proposed DPGNN framework, we select six representative approaches for handling imbalance classification including the current state-of-the-art methods, where the first three target at point-based imbalance classification while the last three are designed specifically for imbalance node classification in networks.

- **Up-sampling:** A classical approach that repeats samples from minority classes. Following [34], we implement this in the embedding space by duplicating minority nodes' representations.
- **Re-weight [33]:** A cost-sensitive approach that assigns class-specific loss weights; we set the weights of each class as inverse ratio of the total training nodes to the number in that class.
- **SMOTE [3]:** Synthetic minority samples are created by interpolating minority samples with their nearest neighbors within the same class based on the output of last GNN layer.
- **GraphSMOTE [34]:** Advancing SMOTE [3], GraphSMOTE has two types of edge generators that can be pre-trained to connect synthetic nodes to the original graph. We report the result of the best generator variant for each dataset.

Table 2: Node classification performance on eight datasets with the best performance emboldened and second underlined.

Model	Dataset (Homophily Value)											
	Cora (0.81)			Citeseer (0.74)			Pubmed (0.80)			Cora-ML (0.79)		
	F1-macro	F1-weight	F1-micro	F1-macro	F1-weight	F1-micro	F1-macro	F1-weight	F1-micro	F1-macro	F1-weight	F1-micro
GCN	0.5205	0.5195	0.5212	0.3870	0.4169	0.4692	0.5501	0.5569	0.5928	0.5205	0.5195	0.5212
GCN _{us}	0.5631	0.5659	0.5727	<u>0.4503</u>	<u>0.4822</u>	<u>0.5220</u>	<u>0.6272</u>	<u>0.6323</u>	<u>0.6451</u>	0.5656	0.5516	0.5611
GCN _{rw}	0.5609	0.5660	0.5724	0.4457	0.4800	0.5156	0.6169	0.6178	0.6327	0.5609	0.5660	0.5724
GCN _{st}	0.5488	0.5398	0.5519	0.4462	0.4794	0.5127	0.5861	0.5964	0.6186	0.5488	0.5398	0.5519
GraphSMOTE	<u>0.5845</u>	<u>0.6026</u>	<u>0.5820</u>	0.4236	0.4774	0.5020	0.6122	0.5998	0.6110	<u>0.6233</u>	<u>0.6450</u>	<u>0.6130</u>
RECT	0.5234	0.5025	0.5448	0.4002	0.4243	0.4549	0.5713	0.5597	0.6002	0.5530	0.5560	0.6026
DR-GCN	0.5513	0.5362	0.5520	0.3924	0.4414	0.4880	0.5628	0.5730	0.5559	0.5412	0.4716	0.5060
DPGCN	0.7115	0.7029	0.7111	0.4838	0.5180	0.5397	0.7018	0.7176	0.7189	0.7273	0.7278	0.7305

Model	Dataset (Homophily Value)											
	DBLP (0.83)			Amazon Computers (0.78)			Amazon Photo (0.83)			Twitch PT (0.58)		
	F1-macro	F1-weight	F1-micro	F1-macro	F1-weight	F1-micro	F1-macro	F1-weight	F1-micro	F1-macro	F1-weight	F1-micro
GCN	0.3482	0.3829	0.3876	0.5343	0.6808	0.6975	0.6999	0.7617	0.7666	0.4557	0.4510	0.4656
GCN _{us}	0.4214	0.4599	0.4795	0.5757	0.6876	0.6883	0.7135	0.7645	0.7632	0.4917	0.5088	0.5131
GCN _{rw}	0.4379	0.4744	<u>0.4892</u>	0.5732	0.6845	0.6841	0.7204	0.7683	0.7670	0.4963	0.5168	0.5193
GCN _{st}	0.3757	0.4302	0.4522	<u>0.5863</u>	0.6999	0.7107	0.7302	0.7782	0.7800	<u>0.5002</u>	0.5267	0.5301
GraphSMOTE	<u>0.4844</u>	<u>0.4938</u>	0.4530	0.5509	0.6213	0.6370	0.7227	0.7716	0.7750	0.3922	0.3558	0.4130
RECT	0.3438	0.3602	0.3810	0.5222	<u>0.7002</u>	0.7351	0.6858	0.7763	<u>0.8007</u>	0.4898	0.4843	0.4904
DR-GCN	0.3797	0.4190	0.4510	0.5357	0.6745	0.7100	<u>0.7434</u>	0.7979	0.8130	0.4791	0.5524	0.5750
DPGCN	0.6167	0.6665	0.6597	0.6702	0.7280	<u>0.7310</u>	0.7600	0.7943	0.7917	0.5600	0.5944	0.5915

**Figure 2: Node classification results under different class imbalance ratios.**

- **RECT** [30]: A supervised model leveraging both a GNN and an unsupervised node proximity-based embedding, which is designed for completely-imbalanced label setting.
- **DR-GCN** [25]: Two types of regularization to tackle class imbalance are proposed: class-conditional adversarial training to separate labeled nodes and unlabeled nodes latent distribution constraint to maintain training equilibrium.

Equipping the first three baselines with a 2-layer GCN [14] encoder, we then have 7 baselines: GCN, GCN_{us}, GCN_{rw}, GCN_{st}, GraphSMOTE, RECT, and DR-GCN that are used for a comprehensive empirical analysis on imbalanced node classification.

5.1.3 Evaluation Metrics. Following existing work in imbalanced classification [9], we measure performance with: F1-macro, F1-micro, and F1-weighted scores. F1-macro and F1-weighted scores evaluate model performance across different classes, with the former taking unweighted mean over the accuracy of each class and the latter weighted mean to account for label imbalance. F1-micro is taken over all testing examples, which gives an overall evaluation of the performance while undervalues nodes in minority classes.

5.1.4 Parameter Settings. We implement our proposed DPGNN and some necessary baselines Pytorch Geometric [22]. For GraphSmote¹, RECT², and DR-GCN³, we use the authors' original code with any needed modifications from their GitHub repositories. Aiming to

provide a rigorous/fair comparison, we tune hyperparameters for all models individually on each dataset around the default/best settings reported in their paper. For DPGCN, we tune the following hyperparameters: dropout rate $\in \{0, 0.5\}$, the coefficient balancing the loss contribution $\lambda_1, \lambda_2 \in \{1, 10\}$, the threshold $\eta \in [0, 6]$, and the training epoch $\in \{1000, 3000, 6000\}$. We select the 2-layer GCN with 256 hidden units as the encoder f_1 due to its simplicity and efficiency, and termed our framework as DPGCN in the following. Note that any other GNN-based encoder can be used instead. For reproducibility, the code of our model is publicly available⁴.

5.2 Performance Comparison

In this subsection, we answer the first question by comparing the performance of DPGCN with other baselines and report the average performance per metric across 20 different data splits along with the edge homophily [36] for each dataset. We observe that DPGCN performs the best in all 8 datasets by F1-macro with a large margin, and best on 7(6) of the 8 datasets according to F1-weight (F1-micro) over other baselines. The F1-micro of GCN is higher than F1-macro on all datasets except Cora, which indicates accuracy is vastly different across different classes and therefore signifies the detrimental effect imposed by class imbalance. GCN_{us, rw, st} improve the performance over GCN on all datasets, which demonstrate their generalizability and effectiveness in handling the imbalance issue. Specifically, GCN_{us} and GCN_{rw} share similar performance since essentially both

¹GraphSMOTE Code: <https://github.com/TianxiangZhao/GraphSmote>

²RECT Code: <https://github.com/zhengwang100/RECT>

³DR-GCN Code: <https://github.com/codeshareabc/DRGCN>

⁴DPGNN Code: <https://github.com/YuWVandy/DPGNN>

of them balance the training loss while the performance of GCN_{st} is unstable since randomly selecting linear interpolation coefficient cannot guarantee the generated instances follow the ground-truth class distributions. GraphSMOTE performs better than other baselines on datasets with higher homophily such as Cora, Cora-ML and DBLP, while worse on datasets with lower homophily such as Citeseer and Twitch PT. This is because edges in lower homophily networks tend to link nodes from different classes and therefore the pre-trained edge generator is also guided to link nodes from different classes. In this case, aggregating neighborhood features incorporates more noise information and therefore compromise the downstream classification. Despite that two advanced models, RECT and DR-GCN, achieve performance gain over the GCN baseline, both of their improvements are surprisingly no obvious compared with other baselines. For example, RECT achieves comparable performance to GCN on Cora and DBLP. For RECT, we hypothesize that the embeddings obtained by GNNs already encode the structural information and therefore further optimizing based on node proximity derives limited useful supervision. Another potential reason is because RECT is proposed for completely-imbalanced label while here we still have some training instances for each minority class. For DR-GCN, the reason of no significant performance improvement is that the limited training nodes may cause over-fitting of their generator and discriminator such that the class distribution is ill-defined via adversarial training [8].

Moreover, we investigate the stability of the performance improvement achieved by our model by varying the imbalance ratio from 1:1 to 1:10. Specifically, we fix the number of training nodes as 2 for minority classes and gradually increasing the number of training nodes from 2 to 40 for majority classes, which exhausts the imbalance scenarios from being balanced to the heavily imbalanced. In Figure 2, we can see DPGCN always performs the best when imbalance ratio is beyond 1:4 on the shown datasets and the gap grows larger as imbalance ratio further increases, which demonstrate the superiority of DPGCN over other baselines in handling heavily imbalanced data. Besides, we observe that our model also achieves higher performance when data is balanced (imbalance ratio 1:1) on most datasets. This is because the imbalanced label propagation derives extra training supervision from pseudo labeled data and indicates the potential of our model in the balanced data setting. On Citeseer, our model achieves lower performance on the first few imbalance ratios because the inefficiency of imbalanced label propagation and smoothing of adjacent nodes due to lower homophily (i.e., 0.74) of Citeseer compared with other four datasets.

5.3 Ablation Study

To demonstrate the contribution of each proposed component in our framework, we conduct ablation studies and answer the second question in this subsection. In Figure 3, we present the performance improvement over GCN achieved by our proposed framework (DPGCN) along with variants that remove the imbalanced label propagation (w/o LP), remove the self-supervised learning components (w/o SSL), and remove both (w/o LP, SSL). We use the best performing hyperparameters found for the results in Table 2 and report the average F1-macro and F1-micro of 20 runs. Note that the most basic model here without SSL and LP has only class prototype-driven balanced training with distance metric learning.

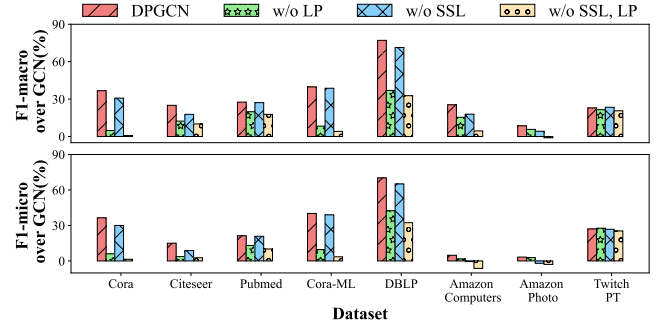


Figure 3: Ablation study results of DPGCN on all datasets.

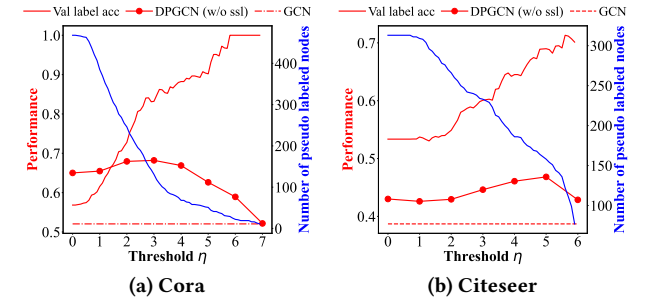


Figure 4: F1-macro score of using different threshold η in imbalanced label propagation on Cora and Citeseer.

Firstly, we observe that every proposed component contributes to the increase of F1-macro, which suggests the validity of each component in tackling class imbalance. However, performance becomes worse than GCN when solely applying class prototype-driven balanced training with distance metric learning (yellow dot bar) on two Amazon Computers and Photos. This is because these two datasets have more classes and therefore the dimension of the concatenated embedding difference g in Eq. (5) becomes pretty high and causes the effect of the curse of high dimensionality [1].

5.4 Influence of Threshold η

Next we answer the last question by analyzing the influence of threshold η in Eq. (12) on the model performance. Here we rule out the SSL component and only consider the performance of DPGCN (w/o SSL), which can exclusively reflect the effect of threshold η in imbalanced label propagation. By varying η from 0 to 7/6 on Cora/Citeseer, we obtain the total number of pseudo labels and their labeling accuracy from Eq. (12). Meanwhile, we visualize the F1-macro score of DPGCN (w/o SSL) and the basic GCN model in Figure 4a. In Cora, as η increases the total number of pseudo labeled nodes decreases while the labeled quality (accuracy) increases since higher η leads to precise labels with sharp distribution. In the initial stage when η is between 0 and 1, the model performance stays the same because even though we obtain extra supervision from additional 500 pseudo labeled nodes as blue curve shows, only around 57% of them are correctly classified and training with these low-quality data would compromise the model performance. As η further rises from 1 to 3, the quality of pseudo labeled nodes improves since their accuracy increases from 0.57 to 0.86 during this stage, thus the F1-macro score of our DPGCN (w/o SSL) rises from 0.66 up to 0.68. After η exceeds 3, even though the quality of the

additional pseudo labels keep increasing, the model performance decreases because we have less pseudo labeled data. Finally, when η is beyond 7 where no additional nodes other than the original imbalanced training nodes are labeled, DPGCN (w/o SSL) degenerate back to the DPGCN (w/o SSL, LP) and the performance gain achieved at this point is mainly due to class prototype-driven balanced training with distance metric learning. In Figure 4b, we observe the similar trend in Citeseer except that in the degenerating stage when η is 6, we still have performance gain over GCN while in Cora, the performance of DPGCN (w/o SSL, LP) converges to GCN, which is consistent with the previous ablation study in Figure 3. Note that here the label accuracy is solely computed for validation nodes by comparing their pseudo labels and ground-truth labels, which can also be treated as a hyperparameter tuning strategy for η .

6 CONCLUSION

In this paper, we focused on the imbalanced node classification problem in graphs, which widely exists in real-world setting such as malicious user detection and drug function testing. Noticing that imbalanced node classification naturally inherits issues of deep learning for imbalance classification: the inclination to learning towards majority classes and the catastrophic forgetting of previous learned instances in minority classes, we use a class prototype-driven balance training scheme to balance the training loss between different classes. To further alleviate the issue of unrepresentative class prototypes, we construct another distance metric space to fully leverage the distance information of query nodes to all class prototypes. Moreover, we derive extra supervision from network topology by an imbalanced label propagation scheme and smooth the learned distance metric representation between adjacent nodes. Experiments on 8 real-world datasets demonstrate the effectiveness of the proposed DPGNN framework in relieving the class imbalance issue. For future work, we plan to study imbalanced graph classification and expect to utilize the graph topology as extra supervision.

REFERENCES

- [1] Charu C Aggarwal, Alexander Hinneburg, and Daniel A Keim. 2001. On the surprising behavior of distance metrics in high dimensional space. In *International conference on database theory*. Springer, 420–434.
- [2] Kelsey Allen, Evan Shelhamer, Hanul Shin, and Joshua Tenenbaum. 2019. Infinite mixture prototypes for few-shot learning. In *International Conference on Machine Learning*. PMLR, 232–241.
- [3] Nitesh V Chawla, Kevin W Bowyer, Lawrence O Hall, and W Philip Kegelmeyer. 2002. SMOTE: synthetic minority over-sampling technique. *Journal of artificial intelligence research* 16 (2002), 321–357.
- [4] Nitesh V Chawla, Nathalie Japkowicz, and Aleksander Kotcz. 2004. Special issue on learning from imbalanced data sets. *ACM SIGKDD explorations newsletter* 6, 1 (2004), 1–6.
- [5] Deli Chen, Yanyai Lin, Lei Li, Xuancheng Ren Li, Jie Zhou, Xu Sun, et al. 2020. Distance-wise Graph Contrastive Learning. *arXiv preprint arXiv:2012.07437* (2020).
- [6] Yu Chen, Lingfei Wu, and Mohammed Zaki. 2020. Iterative deep graph learning for graph neural networks: Better and robust node embeddings. *Advances in Neural Information Processing Systems* 33 (2020).
- [7] Kaize Ding, Jianling Wang, Jundong Li, Kai Shu, Chenghao Liu, and Huan Liu. 2020. Graph prototypical networks for few-shot learning on attributed networks. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*. 295–304.
- [8] Jinhao Dong and Tong Lin. 2019. Margingan: Adversarial training in semi-supervised learning. (2019).
- [9] Yang Gao, Yi-Fan Li, Yu Lin, Charu Aggarwal, and Latifur Khan. 2021. SetConv: A New Approach for Learning from Imbalanced Data. *arXiv preprint arXiv:2104.06313* (2021).
- [10] William L Hamilton. 2020. Graph representation learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning* 14, 3 (2020), 1–159.
- [11] John A Hartigan and Manchek A Wong. 1979. Algorithm AS 136: A k-means clustering algorithm. *Journal of the royal statistical society. series c (applied statistics)* 28, 1 (1979), 100–108.
- [12] Haibo He and Yunqian Ma. 2013. Imbalanced learning: foundations, algorithms, and applications. (2013).
- [13] Justin M Johnson and Taghi M Khoshgoftaar. 2019. Survey on deep learning with class imbalance. *Journal of Big Data* 6, 1 (2019), 1–54.
- [14] Thomas N. Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *5th International Conference on Learning Representations, ICLR*.
- [15] Johannes Klicpera, Aleksandar Bojchevski, and Stephan Günnemann. 2019. Predict then Propagate: Graph Neural Networks meet Personalized PageRank. In *7th International Conference on Learning Representations, ICLR*.
- [16] Bartosz Krawczyk. 2016. Learning from imbalanced data: open challenges and future directions. *Progress in Artificial Intelligence* 5, 4 (2016), 221–232.
- [17] Qimai Li, Zhichao Han, and Xiao-Ming Wu. 2018. Deeper Insights Into Graph Convolutional Networks for Semi-Supervised Learning. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*.
- [18] Meng Liu, Hongyang Gao, and Shuiwang Ji. 2020. Towards Deeper Graph Neural Networks. In *KDD '20: The 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Virtual Event, CA, USA, August 23–27, 2020*.
- [19] Wei Liu, Junfeng He, and Shih-Fu Chang. 2010. Large graph construction for scalable semi-supervised learning. In *ICML*.
- [20] Yixin Liu, Shirui Pan, Ming Jin, Chuan Zhou, Feng Xia, and Philip S Yu. 2021. Graph self-supervised learning: A survey. *arXiv preprint arXiv:2103.00111* (2021).
- [21] 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.
- [22] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *arXiv preprint arXiv:1912.01703* (2019).
- [23] Benedek Rozemberczki, Carl Allen, and Rik Sarkar. 2021. Multi-scale attributed node embedding. *Journal of Complex Networks* 9, 2 (2021), cnab014.
- [24] Oleksandr Shchur, Maximilian Mumme, Aleksandar Bojchevski, and Stephan Günnemann. 2018. Pitfalls of graph neural network evaluation. *arXiv preprint arXiv:1811.05868* (2018).
- [25] Min Shi, Yufei Tang, Xingquan Zhu, David Wilson, and Jianxun Liu. 2020. Multi-Class Imbalanced Graph Convolutional Network Learning. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence (IJCAI-20)*.
- [26] Jake Snell, Kevin Swersky, and Richard S. Zemel. 2017. Prototypical Networks for Few-shot Learning. In *Advances in Neural Information Processing Systems 30 (NeurIPS)*. 4077–4087.
- [27] Nguyen Thai-Nghe, Zeno Gantner, and Lars Schmidt-Thieme. 2010. Cost-sensitive learning methods for imbalanced data. In *The 2010 International joint conference on neural networks (IJCNN)*. IEEE, 1–8.
- [28] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph Attention Networks. In *International Conference on Learning Representations*.
- [29] Oriol Vinyals, Charles Blundell, Tim Lillicrap, Koray Kavukcuoglu, and Daan Wierstra. 2016. Matching Networks for One Shot Learning. In *Advances in Neural Information Processing Systems 29 (NeurIPS)*. 3630–3638.
- [30] Zheng Wang, Xiaojun Ye, Chaokun Wang, Jian Cui, and Philip Yu. 2020. Network embedding with completely-imbalanced labels. *IEEE Transactions on Knowledge and Data Engineering* (2020).
- [31] Lingfei Wu, Ian En-Hsu Yen, Zhen Zhang, Kun Xu, Liang Zhao, Xi Peng, Yinglong Xia, and Charu Aggarwal. 2019. Scalable global alignment graph kernel using random features: From node embedding to graph embedding. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1418–1428.
- [32] Zhilin Yang, William Cohen, and Ruslan Salakhudinov. 2016. Revisiting semi-supervised learning with graph embeddings. In *International conference on machine learning*. PMLR, 40–48.
- [33] Bo Yuan and Xiaoli Ma. 2012. Sampling+ reweighting: Boosting the performance of AdaBoost on imbalanced datasets. In *The 2012 international joint conference on neural networks (IJCNN)*. IEEE, 1–6.
- [34] Tianxiang Zhao, Xiang Zhang, and Suhang Wang. 2021. GraphSMOTE: Imbalanced Node Classification on Graphs with Graph Neural Networks. In *Proceedings of the ACM International Conference on Web Search and Data Mining*. 833–841.
- [35] Fan Zhou, Chengtai Cao, Kunpeng Zhang, Goce Trajcevski, Ting Zhong, and Ji Geng. 2019. Meta-gnn: On few-shot node classification in graph meta-learning. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*. 2357–2360.
- [36] Jiong Zhu, Yujun Yan, Lingxiao Zhao, Mark Heimann, Leman Akoglu, and Danai Koutra. 2020. Beyond Homophily in Graph Neural Networks: Current Limitations and Effective Designs. *Advances in Neural Information Processing Systems* 33.