Ripple Walk Training: A Subgraph-based Training Framework for Large and Deep Graph Neural Network

Jiyang Bai[§], Yuxiang Ren[§] and Jiawei Zhang
Department of Computer Science, Florida State University, FL, USA
Email: bai@cs.fsu.edu, yuxiang@ifmlab.org, jiawei@ifmlab.org

***Abstract*—Graph neural networks (GNNs) have achieved outstanding performance in learning graph-structured data and various tasks. However, many current GNNs suffer from three common problems when facing large-size graphs or using a deeper structure: *neighbors explosion*, *node dependence*, and *oversmoothing*. Such problems attribute to the data structures of the graph itself or the designing of the multi-layers GNNs framework, and can lead to low training efficiency and high space complexity. To deal with these problems, in this paper, we propose a general subgraph-based training framework, namely Ripple Walk Training (RWT), for deep and large graph neural networks. RWT samples subgraphs from the full graph to constitute a mini-batch, and the full GNN is updated based on the mini-batch gradient. We analyze the high-quality subgraphs required in a mini-batch in a theoretical way. A novel sampling method Ripple Walk Sampler works for sampling these high-quality subgraphs to constitute the mini-batch, which considers both the randomness and connectivity of the graph-structured data. Extensive experiments on different sizes of graphs demonstrate the effectiveness and efficiency of RWT in training various GNNs (GCN & GAT). Our code is released in the https://github.com/anonymous2review/RippleWalk.**

I. Introduction

Graph neural networks (GNNs) have achieved outstanding performance in graph-structured data based applications, such as knowledge graphs [28], social medias [22], and protein interface prediction [6]. GNNs learn nodes' high-level representations through a recursive neighborhood aggregation scheme [29]. As the graph's scale increases and higher-order neighbors are considered, the recursive neighborhood aggregation can cause the number of neighbors to explode. We name this problem as *neighbors explosion*. For each node, the *neighbors explosion* will lead to the computation complexity exponentially increasing with the GNNs depth [5] and the graph size. Therefore, some current works on GNNs (e.g., graph attention networks (GAT) [27]) can only handle small-size graphs (normally less than 5000 nodes) with a shallow structure (less than 3 layers). Simultaneously, the graph-structured data has the characteristics of *node dependence*, which means neighboring nodes affect each other in the learning process. As a result, the most current GNNs have to learn on the full graph, and when the size of the graph is too large, it is easy to reach the upper limit of the device memory. Even with a shallow GNNs structure, the memory space demand is extremely large. *node dependence* also limits the performance of training methods such as mini-batch SGD. Because calculating the loss of one node, GNNs require embeddings of all node neighbors, and its neighbors also need embeddings of their neighbors for aggregation. This increases the overhead of mini-batch SGD, especially for dense graphs and deeper GNNs. Another factor that limits the effectiveness of GNNs is *oversmoothing*. Especially when GNNs go deeper and learn on the full graph, and it is unavoidable that node representations from different clusters mix up [33]. But this aggregation is unexpected because nodes from different clusters do not meet the smoothness assumptions on the graph (close nodes are similar). Finally, the *oversmoothing* will lead to node representations indistinguishable. Therefore, when the GNN has a deep structure, not only the training is more difficult, but *oversmoothing* also impedes its performance.

To deal with the three problems mentioned above, some methods have emerged. GraphSAGE [8] learns a function that generates embeddings by sampling and aggregating features from a node's local neighborhood. FastGCN [4] utilizes Monte Carlo approaches to sample neighbors which avoids the *neighbors explosion*. Chen et.al. [3] develop control variate based algorithms that allow sampling an arbitrarily small neighbor size. They all use neighbor sampling to avoid *neighbors explosion* and improve the training speed, but they can not handle the remaining problems. When the full graph's size is large, the memory overhead for learning on the full graph is unacceptable. These methods do not optimize the memory overhead when speed up training. Cluster-GCN [5] has a training algorithm based on subgraphs, which are constructed by clustering on the full graph. The subgraphs are selected randomly to constitute mini-batches to train the GCN. However, the size of clusters in a graph is difficult to control. When very large subgraphs are constructed based on the clustering results, Cluster-GCN lacks scalability and cannot tackle the *neighbors explosion*. Besides, the time and space overhead of clustering on a large graph are also nonnegligible.

In this paper, we propose a general subgraph-based training framework, namely **R**ipple **W**alk **T**raining (RWT), for deep and large graph neural networks. RWT aims to handle all aforementioned problems simultaneously. RWT is developed from the mini-batch training, but there exist apparent differences. Instead of sampling neighbors and training on the

[§]Equal contribution

full graph, RWT samples subgraphs from the full graph to constitute a mini-batch. The full GNN is updated based on the mini-batch gradient. We design a novel sampling method Ripple Walk Sampler for RWT, which considers both the randomness and connectivity of the graph-structured data to deal with those three problems. RWT can sample high-quality subgraphs to constitute the mini-batch to benefit efficient training. For the problem of *neighbors explosion*, the mini-batch gradient is calculated within subgraphs so that subgraphs of acceptable size can completely avoid this problem. At the same time, the gradient does not depend on nodes outside the subgraph, which solves the *node dependence* at the subgraph level. Unexpected aggregations usually occur between subgraphs. Yet, the propagation-aggregation happens within the subgraph, so the *oversmoothing* can be handled.

The contributions of our work are summarized as follows:

- We propose a general subgraph-based training framework Ripple Walk Training (RWT) for GNNs. RWT not only accelerates the training speed on the large graph but also breaks through the memory bottleneck. In addition, it can effectively deal with the problem of the *oversmoothing* that occurs in deep GNNs.
- We analyze what kind of subgraphs can support effective and efficient training. Based on the analysis, we design a novel sampling method Ripple Walk Sampler with the theoretical guarantee.
- We conduct extensive experiments on different sizes of graphs to demonstrate the effectiveness of RWT. The results show the superiority of Ripple Walk in training different GNNs (GCN & GAT) subject to the performance and training efficiency.

II. Related Works

Graph neural networks (GNNs) aim at the machine learning tasks involving graph-structured data. The first research work extending the convolutional neural network to the graph-structured data is [26]. After that, more related works [24], [19] were introduced. More recently, [2] is proposed and based on spectral graph theory. Later, spatial-based ConvGNNs [27], [20] define graph convolutions directly based on a node's spatial relations. The spectral-based and spatial-based GNNs can be regarded as an information propagation-aggregation mechanism, and such mechanism is achieved by the connections and multi-layer structure. Although GNNs have outstanding performance, they are also limited by the problems from three aspects: *node dependence*, *neighbors explosion*, and *oversmoothing*. Aiming at these problems, some related works have been proposed in different directions.

Node dependence [5] forces GNNs to be trained on the entire graph, which leads to the slow training process. More specifically, in each training epoch, the information aggregation involves the full graph's adjacency matrix. Thus the space complexity will be at least $O(|\mathcal{V}|^2)$, where $|\mathcal{V}|$ is the size of a full graph. To deal with such a problem, [5], [31] apply the concept of subgraph training methods. The essence of subgraph training is to collect a batch of subgraphs from the full graph and use them during the training process. There are also other approaches to optimize the GNNs frameworks. [12], [13], [16] optimize the localized filter to reduce the time cost of training on the full graph. Further, [9], [15] reduce the number of learnable parameters by dimensionality reduction and residual graph Laplacian, respectively. But these approaches do not alleviate the space complexity problem.

Neighbors expansion makes deep GNNs difficult to being implemented. Because learning a single node requires embeddings from its neighbors, and the quantity may be explosive when a GNN goes deeper. Some research works deal with *neighbors explosion* by neighbors sampling [8], [4], [3]. In other directions, several models [7], [30] select specific neighbors based on defined metrics to avoid the explosive quantity. [23] randomly remove edges from input graphs to handle the *neighbor explosion*. The works mentioned above all focus on the neighbor-level sampling but still have the same space complexity with original GNNs.

The problem of *oversmoothing* in the GCN was introduced in [14]. When GNNs go deep, the performance suffers from *oversmoothing*, where node representations from different clusters become mixed up [33]. The node information, or representation propagation-aggregation mechanism, can be regarded as one type of random walk within the graph. With the increasing of walking steps, the representations of nodes will finally converge to a stable status. Such convergence would impede the performance of GNNs and make the nodes indistinguishable in the downstream tasks. Some related works have been proposed to deal with the *oversmmothing*. [32] comes up with the suspended animation and utilizes the residual networks to mine the advantages of deeper networks.

III. Proposed Algorithm

A. Preliminaries and Background

For most widely used GNNs models (e.g., GCN, GAT), the essence of which is aggregating feature representation for each node in the full graph, and then using the aggregated feature representation to accomplish specific tasks. Given a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, the aggregation procedure of GNNs layers is shown as follow:

$$\begin{aligned} &\mathbf{h}^{(0)} = \mathbf{X} \\ &\mathbf{h}^{(l+1)}[i] = \sigma(\sum_{j \in \mathcal{N}_i} \alpha_{ij} \cdot \mathbf{h}^{(l)}[j]\mathbf{W}^{(l)}) \end{aligned} \tag{1}$$

Here, the $\mathbf{X} \in \mathbb{R}^{|\mathcal{V}| \times F_0}$ is the input feature vectors (matrix) of all the nodes in graph $\mathcal{G}$; $\mathbf{h}^{(l)}[i]$ is the hidden feature of node i in the l_{th} layer; σ is the non-linear function such as Relu [21]; $\mathbf{W}^{(l)}$ is the learnable linear transfer matrix; α is a variant of adjacency matrix, which represents different meanings according to different GNNs models. For example, in GCN structure, $\alpha = \widetilde{\mathbf{A}}$ is the normalized adjacency matrix. During the feedforward process, the hidden representations of node i are updated by aggregating both its features and the local neighbors' hidden features. After layers of computing, the output representations of nodes will be delivered to the

Algorithm 1 Ripple Walk Training for GNNs

Input: Graph $\mathcal{G}$; GNNs model $H_{\mathbf{W}}(\cdot)$; loss function $Loss(\cdot)$; training iteration number T; subgraph mini-batch size M
Output: Trained $H_{\mathbf{W}}(\cdot)$

```
Initialize subgraph mini-batch batch = {}
for k = 1, 2, ..., M do
    G_k ← Ripple Walk Sampler /* By Algorithm 2 */
    batch = batch ∪ {G_k}
end for
for t = 1, 2, ..., T do
    Select a subgraph from batch as G_t
    loss = Loss(H_W(G_t), y_{G_t})  /* The y_{G_t} denotes the ground truth of nodes in G_t. */
    Update W according to the gradient ∇_W loss
end for
return H_W(·)
```

downstream tasks. The learnable weights will be optimized during the backpropagation.

The calculation in Equation 1 uses the full graph of $\mathcal{G}$, and the full graph involved will easily lead to the concern of *node dependence* and *neighbors explosion*. Since it requires the full adjacency matrix $\mathbf{A}$ and entire feature matrix $\mathbf{X}$, the size of which would be too large to deal with. Both the increasing size of the graph (e.g., graph with millions of nodes [5]) and more sophisticated models (e.g., deeper layers [32]) would aggravate the problems.

B. Subgraph-based Training

To solve the problem of computationally expensive, an alternative approach is training the GNNs with RWT. The procedure of RWT is presented in Algorithm 1. In Algorithm 1, the subgraph mini-batch size M are varying for different datasets. Generally, the value of M is to satisfy that the overall number of sampled nodes (in all subgraphs) is about ten times the full graph size. Unlike the training process involving the full graph $\mathcal{G}$, RWT employs a subgraph of $\mathcal{G}$ in each training iteration. In other words, a smaller size of α matrix and only part of the nodes are required during each training epoch. In this way, the aggregation procedure in the t_{th} training iteration is

$$\begin{aligned}\mathbf{h}^{(0)} &= \mathbf{X}_{\mathcal{G}_t} \\ \mathbf{h}^{(l+1)}[i] &= \sigma(\sum_{j\in\mathcal{N}_i^t} \alpha_{ij}^t \cdot \mathbf{h}^{(l)}[j]\mathbf{W})\end{aligned} \tag{2}$$

Here, the $\mathcal{G}_t = (\mathcal{V}_t, \mathcal{E}_t)$ is a subgraph of $\mathcal{G}$, where $\mathcal{V}_t \subseteq \mathcal{V}$ and $\mathcal{E}_t \subseteq \mathcal{E}$; $\mathcal{N}_i^t$ is the neighbor nodes set of node i in $\mathcal{G}^t$; α^t corresponds to the adjacency matrix of $\mathcal{G}_t$. For different training iterations, different subgraphs will be employed into Equation 2. Comparing to the Equation 1, the computational complexity in Equation 2 can be reduced from $O(|\mathcal{N}||\mathcal{V}|)$ to $O(|\mathcal{N}^t||\mathcal{V}_t|)$.

The switch from Equation 1 to Equation 2 is similar to the change from gradient descent to mini-batch gradient descent. For RWT, the concerns are also reflected in two aspects: (1) each subgraph only contains part of the nodes; (2) subgraph is equivalent to dropping some edges, which means the dependency (connections) of nodes is incomplete. Unlike previous data type (e.g., image data), where each data sample is independent, the graph type data consists of tons of nodes connected to others. To respond to these concerns and prove the effectiveness of GNNs models with subgraphs, we propose the following theorems.

THEOREM 1. *Given graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, assume the $\mathcal{V}' \subseteq \mathcal{V}$ and the nodes in $\mathcal{V}'$ are randomly sampled from $\mathcal{V}$; H is a GNNs structure. The objective fucntion of training H with subset nodes ($\mathcal{V}'$) (with all neighbors) is equivalent to the objective fucntion of training with full graph, which can be presented as:*

$$\min_H \frac{1}{|\mathcal{V}|}\sum_{i\in\mathcal{V}} loss(H(\mathcal{G}(i)), \mathbf{y}_i) \doteq \min_H \frac{1}{|\mathcal{V}'|}\sum_{j\in\mathcal{V}'} loss(H(\mathcal{G}(j)), \mathbf{y}_j) \tag{3}$$

where $\doteq$ denotes unbiased estimation; $loss(\cdot)$ is the selected loss function; $\mathcal{G}(i)$ means using node i's neighbors in $\mathcal{G}$ (all neighbors) during neighbors aggregation.

Proof. Similar to the switch from gradient descent (GD) to stochastic gradient descent (SGD), where the gradient calculated in SGD is an estimation of that in GD, the proof of Theorem 1 also follows the same rule. For the loss fucntion with full graph $\mathcal{G}$ (left part of Equation 3), it has

$$\begin{aligned}\frac{1}{|\mathcal{V}|}\sum_{i\in\mathcal{V}} loss(H(\mathcal{G}(i)), \mathbf{y}_i) &= \frac{1}{|\mathcal{V}|}|\mathcal{V}| \cdot \mathbb{E}_{i\in\mathcal{V}}[loss(H(\mathcal{G}(i)), \mathbf{y}_i)] \\ &= \mathbb{E}_{i\in\mathcal{V}}[loss(H(\mathcal{G}(i)), \mathbf{y}_i)]\end{aligned} \tag{4}$$

In the above equation, the loss function is expressed as the expectation format. Let us denote the loss function with $\mathcal{V}'$ (right part of Equation 3) as $\mathcal{L}'$, $\mathcal{L}' = \frac{1}{|\mathcal{V}'|}\sum_{j\in\mathcal{V}'} loss(H(\mathcal{G}(j)), \mathbf{y}_j)$. Since the nodes in $\mathcal{V}'$ are randomly sampled from $\mathcal{V}$, according to the statistical leanrning [1] the $\mathcal{L}'$ is an unbiased estimation of $\mathbb{E}_{i\in\mathcal{V}}[loss(H(\mathcal{G}(i)), \mathbf{y}_i)]$. Thus, we have $\mathbb{E}_{i\in\mathcal{V}}[loss(H(\mathcal{G}(i)), \mathbf{y}_i)] \doteq \mathcal{L}'$, which also means

$$\min_H \frac{1}{|\mathcal{V}|}\sum_{j\in\mathcal{V}} loss(H(\mathcal{G}_(j)), \mathbf{y}_j) \doteq \min_H \frac{1}{|\mathcal{V}'|}\sum_{i\in\mathcal{V}'} loss(H(\mathcal{G}(i)), \mathbf{y}_i)$$

□

THEOREM 2. *Under the settings in Theorem 1, the objective function of training H with subgraph $\mathcal{G}' = (\mathcal{V}', \mathcal{E}')$ is equivalent to training with subset nodes $\mathcal{V}'$ (with all neighbors), which can be represented as*

$$\min_H \frac{1}{|\mathcal{V}'|}\sum_{i\in\mathcal{V}'} loss(H(\mathcal{G}(i)), \mathbf{y}_i) \doteq \min_H \frac{1}{|\mathcal{V}'|}\sum_{j\in\mathcal{V}'} loss(H(\mathcal{G}'(j)), \mathbf{y}_j) \tag{5}$$

where $\mathcal{G}'(j)$ means using node j's neighbor in $\mathcal{G}'$ (partial neighbors).

Proof. The only difference between these two objective functions is the neighbors of each node. Since only the subgraph $\mathcal{G}' = (\mathcal{V}', \mathcal{E}')$ are involved during the training process, for node $i \in \mathcal{V}'$, only part of its neighbors are also in $\mathcal{V}'$. In other words, $\mathcal{N}_i' \subseteq \mathcal{N}_i$, where $\mathcal{N}_i'$ is the neighbor set of node i in subgraph $\mathcal{G}'$. According to [10], the feed-forward propagation of node i can be expressed as

$$\mathbf{h}^{(l+1)}[i] = \sigma(\sum_{k\in\mathcal{N}_i} \mathbf{h}^{(l)}[k] \cdot \mathbf{W}^{(l)}) \doteq \sigma(|\mathcal{N}_i| \cdot \mathbb{E}_{k\in\mathcal{N}_i}[\mathbf{h}^{(l)}[k]] \cdot \mathbf{W}^{(l)}) \tag{6}$$

Algorithm 2 Ripple Walk Sampler

Input: Target graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$; expansion ratio r; target subgraph size S
Output: Subgraph $\mathcal{G}_k$

```
Initiate G_k = (V_k, E_k) with V_k = φ
Randomly select the initial node v_s, add v_s into the G_k
while |V_k| < S do
    NS = {n|(n, j) ∈ E, j ∈ V_k, n ∈ V \ V_k} /* Get neighbor nodes set of current V_k */
    Randomly select r of nodes in NS, add them into the V_k
end while
return G_k
```

The expectation $\mathbb{E}_{k\in\mathcal{N}_i}[\mathbf{h}^{(l)}[k]]$ in the above equation can be estimated by the

$$\mathbb{E}_{k\in\mathcal{N}_i}[\mathbf{h}^{(l)}[k]] \doteq \frac{1}{|\mathcal{N}_i'|}\sum_{k\in\mathcal{N}_i'}\mathbf{h}^{(l)}[k] \tag{7}$$

if the nodes in $\mathcal{N}_i'$ are randomly selected from $\mathcal{N}_i$. Given node $k \in \mathcal{N}_i$, we denote the possibility that node k will be selected into $\mathcal{V}'$ as $p(k|i)$. We know that $\forall k, h \in \mathcal{N}_i$, $\mathrm{p}(k|i) = \mathrm{p}(k) = \mathrm{p}(h) = \mathrm{p}(h|i)$ in every step. Thus the Equation 7 can be satisfied, and

$$H(\mathcal{G}(i)) \doteq H(\mathcal{G}'(i)),\ \forall i \in \mathcal{V}' \tag{8}$$

Therefore, the Equation 5 can hold. □

From the analysis above, to achieve the equivalent training effect, the subgraphs should possess:

- **randomness**: The randomness contains two aspects: (1) each node has the same probability to be selected; (2) for any node, its neighbors own the same probability to be selected. Randomness can help eliminate the *neighbors explosion* problem.
- **connectivity**: The subgraph should preserve the connectivity in the full graph. The connectivity of each subgraph should be high enough to preserve the connectivity in the full graph. This corresponds to the *node dependence* problem.

In this way, even though each subgraph cannot singly cover all the nodes and structure information in $\mathcal{G}$, the batch of subgraphs can help achieve the same object as the full graph as long as each subgraph satisfies the randomness and connectivity characteristics. To follow these two characteristics, we propose the Ripple Walk Sampler algorithm.

C. Ripple Walk Subgraph Sampling

The Ripple Walk Sampler algorithm is shown in Figure 1. For the subgraph $\mathcal{G}_k$, it is initialized with a random node v_s, then expands along with the connections among nodes. After multiple expansion steps (sampling), the subgraph with a specific size (e.g., S) will be returned. During each expansion step, the neighbor set (shown by the background color region in Figure 1) contains the potential nodes to be sampled. Then r (e.g., $r = 0.5$) of the nodes (shown by colored nodes in Figure 1) in neighbor set will be added into the current subgraph. Here, r is the expansion ratio, which means the ratio of nodes in the neighbor set to be sampled in the current step. Such an expansion process operates like the "ripple" on the water. More details of Ripple Walk Sampler are exhibited in Algorithm 2.

From the analysis in Section III-B, we conclude that it is ideal if the sampled subgraphs possess both randomness and connectivity. The Ripple Walk Sampler strategy can maintain randomness by randomly sampling initial node and random expansion, while the expansion along edges can guarantee connectivity. In the following part, we will show the advantages of Ripple Walk Sampler algorithm concerning those two characteristics.

During the expansion of Ripple Walk Sampler, r determines the range of the subgraph. When $r \to 0$, it can be regarded as random sampling. For the randomly sampled subgraph, the connectivity might be too low to reproduce the global structure in the full graph. To show the advantages of Ripple Walk Sampler compared with a random sample, we first state the following theorem:

THEOREM 3. *From graph $\mathcal{G}$, $\mathcal{G}_k = (\mathcal{V}_k, \mathcal{E}_k)$ is the subgraph generated by Ripple Walk Sampler, while $\mathcal{G}_r = (\mathcal{V}_r, \mathcal{E}_r)$ is the randomly sampled subgraph. Then $\forall i, j \in \mathcal{V}_r$ and $\forall m, l \in \mathcal{V}_k$, it has*

$$p((i,j) \in \mathcal{E}_r) \le p((m,l) \in \mathcal{E}_k) \tag{9}$$

where $p(\cdot)$ denotes the probability.

Proof. According to Algorithm 2, in each sampling step, for $\forall i \in NS$, there $\exists j \in \mathcal{V}_k$ having $(i,j) \in \mathcal{E}$. Thus when Ripple Walk adds one node into the subgraph, one edge will be added into $\mathcal{E}_k$ as well. For $\mathcal{G}_r$, when a new node is selected into the subgraph, possibly there is no new edge added. For subgraphs with the same number of nodes, more connections will selected into $\mathcal{E}_k$ comparing to $\mathcal{E}_r$. Therefore, we have $\mathrm{p}((i,j) \in \mathcal{E}_r) \le \mathrm{p}((m,l) \in \mathcal{E}_k)$. □

From Theorem 3, it is obvious that Ripple Walk can join more connections during the sampling process. Thus the connectivity of subgraphs by Ripple Walk Sampler is higher than the randomly sampled subgraphs.

Similar to the Ripple Walk, Breadth-First-Search (BFS) is a graph search algorithm that expands from one central node and traverses the whole neighbor set. Essentially, BFS is equivalent to Ripple Walk Sampler with $r \to 1$. Different from Ripple Walk Sampler, BFS cannot guarantee the randomness of node sampling: for BFS, once the initial node v_s and the target subgraph size S are certain, the nodes to be selected into the subgraph have been determined. In fact, if BFS satisfies the randomness mentioned in Section III-B, the subgraph cannot be determined by the initial node. On the other hand, Ripple Walk Sampler can maintain randomness. Except for the random initial node, the neighbor nodes in each step are sampled randomly. Even starting from the same initial node, Ripple Walk Sampler can still generate different subgraphs.

From the analysis above, Ripple Walk Sampler not only can keep the randomness of the sampled subgraphs, but also maintain a relatively high level of connectivity. With these two characteristics, RWT can solve the *neighbors explosion* and *node dependence* problems, meanwhile reproduce the information in the full graph. The selection of expansion ratio

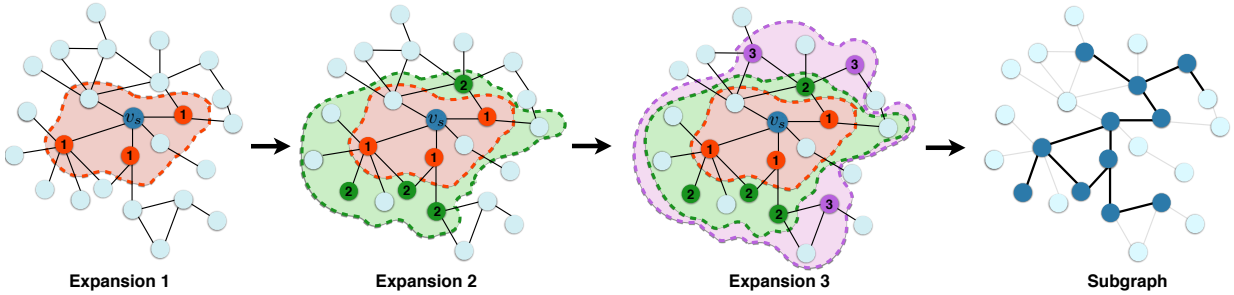

Fig. 1. Ripple Walk Sampling. In each expansion, the expansion ratio $r = 0.5$, the background color region represents the neighbor set, the colored nodes represents the truly sampled nodes. (Best viewed in color)

is important, and we provide the analyses of parameter r in Section IV-C1.

D. For Deeper Graph Networks

The commonly used GNNs only involve no more than two layers. According to [14], each GCN layer can be regarded as one type of Laplacian smoothing, which essentially computes the features of nodes as the weighted average of itself and its neighbors'. In other words, GNN structures with much deeper layers will repeatedly carry out Laplacian smoothing, and features of nodes will finally converge to the global steady states. Such smoothing will undermine the learning ability of GNNs. This point of view also corresponds to the concepts of over smoothing and mixing time in [23], [17].

By applying GNNs with subgraphs, we will prove that RWT can eliminate the problem of converging to global steady states. Subsequently, GNNs with deeper layers can achieve better learning capability. We will give the following definition and assumption.

Definition 1. *(Node distribution): In graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, $\mathbf{h}^{(0)}[i] \sim \mathcal{D}_i$ denotes that the feature representation of node $i \in \mathcal{V}$ is under the distribution $\mathcal{D}_i$.*

In graph $\mathcal{G}$, each node is under a corresponding distribution. While different nodes might own different labels, we assume that nodes within the same class share similar distributions. The assumption can be presented as

Proposition 1. *In graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with $i, j, k \in \mathcal{V}$, if $\mathbf{y}_i = \mathbf{y}_j$ and $\mathbf{y}_i \neq \mathbf{y}_k$, then we assume*

$$\mathbb{E}_{\mathbf{y}_i=\mathbf{y}_j}[KL(\mathcal{D}_i, \mathcal{D}_j)] \leq \mathbb{E}_{\mathbf{y}_i\neq\mathbf{y}_k}[KL(\mathcal{D}_i, \mathcal{D}_k)] \tag{10}$$

where $\mathbb{E}$ is the expectation and KL is the Kullback-Leibler divergence (KL divergence).

KL divergence is a measure of the difference between two probability distributions. To be simplified, here we call it the KL divergence of two nodes. It is easy to understand since the same labeled nodes are more likely to share information (features) that comes from similar distributions.

According to Equation 1, the computation in each GNNs layer is the weighted averaging of each node's neighbors. If we ignore the linear transform by $\mathbf{W}^{(0)}$, from the node distribution view, it can be written as

$$\mathbf{h}^{(1)}[i] = \sum_{k\in\mathcal{N}_i} \alpha_{ik}\mathbf{h}^{(0)}[k] \sim \text{Joint}(\mathcal{D}_{k\in\mathcal{N}_i}) \triangleq \mathcal{D}_i^{(1)} \tag{11}$$

where $\text{Joint}(\mathcal{D}_{k\in\mathcal{N}_i})$ means the weighted average distribution of each $\mathcal{D}_k$, and we denote $\text{Joint}(\mathcal{D}_{k\in\mathcal{N}_i})$ as $\mathcal{D}_i^{(1)}$. Through one layer of calculation, the new hidden representation of node i will be under the $\text{Joint}(\mathcal{D}_{k\in\mathcal{N}_i})$ distribution. After l layers, we denote it as $\mathbf{h}^{(l)}[i] \sim \mathcal{D}_i^{(l)}$.

Theorem 4. *For full graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ and subgraph mini-batch $\{\mathcal{G}_1, \mathcal{G}_2, \ldots, \mathcal{G}_M\}$ generated by Ripple Walk Sampler. Assume the nodes within the local parts are more likely to share the same label. Let $i, j \in \mathcal{V}$ and $m, n \in \mathcal{V}_k$, $\mathcal{G}_k = (\mathcal{V}_k, \mathcal{E}_k) \in \{\mathcal{G}_1, \mathcal{G}_2, \ldots, \mathcal{G}_M\}$. Then,*

$$\mathbb{E}_{m,n\in\mathcal{V}_k}[KL(\mathcal{D}_m, \mathcal{D}_n)] \leq \mathbb{E}_{i,j\in\mathcal{V}}[KL(\mathcal{D}_i, \mathcal{D}_j)] \tag{12}$$

Proof. According to Ripple Walk Sampler, $\mathcal{G}_k$ only covers part of local nodes in $\mathcal{G}$. Thus for $\forall i, j \in \mathcal{V}$ and $\forall m, n \in \mathcal{V}_k$, $\text{p}(\mathbf{y}_i = \mathbf{y}_j) \leq \text{p}(\mathbf{y}_m = \mathbf{y}_n)$. Therefore,

$$\begin{aligned}
&\mathbb{E}_{m,n\in\mathcal{V}_k}[KL(\mathcal{D}_m, \mathcal{D}_n)] \\
=&\text{p}(\mathbf{y}_m = \mathbf{y}_n) \cdot \mathbb{E}_{\mathbf{y}_m=\mathbf{y}_n} KL(\mathcal{D}_m, \mathcal{D}_n) \\
&+ \text{p}(\mathbf{y}_m \neq \mathbf{y}_n) \cdot \mathbb{E}_{\mathbf{y}_m\neq\mathbf{y}_n} KL(\mathcal{D}_m, \mathcal{D}_n) \\
\leq&\text{p}(\mathbf{y}_i = \mathbf{y}_j) \cdot \mathbb{E}_{\mathbf{y}_i=\mathbf{y}_j} KL(\mathcal{D}_i, \mathcal{D}_j) \\
&+ \text{p}(\mathbf{y}_i \neq \mathbf{y}_j) \cdot \mathbb{E}_{\mathbf{y}_i\neq\mathbf{y}_j} KL(\mathcal{D}_i, \mathcal{D}_j) \\
=&\mathbb{E}_{i,j\in\mathcal{V}}[KL(\mathcal{D}_i, \mathcal{D}_j)]
\end{aligned} \tag{13}$$

□

From Theorem 4, the distribution similarity of nodes in the subgraph is higher than that in the full graph. It is easy to understand since each subgraph generated by Ripple Walk contains nodes from the local part. The randomness ensures that different subgraphs could cover other local parts of the full graph for the subgraph mini-batch. Subsequently, with the increasing of l, the distribution $\mathcal{D}^{(l)}$ in each subgraph will converge to different steady states: since each subgraph possesses different nodes and structures. Compared with the global steady state, different steady states correspond to the local information within various subgraphs, which can help improve the learning capacity of deep GNNs.

IV. Experiments

To show the effectiveness and efficiency of RWT, extensive experiments have been conducted on real-world datasets. We aim to answer the following evaluation questions based on experimental results together with the detailed analysis:

- **Question 1**: Can RWT break through the memory bottleneck in order to handle *node dependence*?
- **Question 2**: Can RWT accelerate the training speed on large graphs when facing *neighbors explosion*?

TABLE I
DATASETS IN EXPERIMENTS

	Transductive			Inductive	
	Cora	Citeseer	Pubmed	Flickr	Reddit
# Nodes	2708	3327	19717	89250	232965
# Edges	5429	4732	44338	899756	11606919
# Features	1433	3703	500	500	602
# classes	7	6	3	7	41
Label rate	0.052	0.036	0.003	0.6	0.6

- **Question 3**: Can Ripple Walk Sampler provide powerful subgraphs to support effective and efficient training?
- **Question 4**: Can RWT tackle the *oversmoothing* problem occuring in deep GNNs?

A. Experiment Settings

1) Datasets: We test our algorithms on 5 datasets. Three of them are standard citation network benchmark datasets: Cora, Citeseer, and Pubmed [25]. Flickr [31], [18] is built by forming links between images sharing common metadata from Flickr. Edges are formed between images from the same location, submitted to the same gallery, group, or set, images sharing common tags, images taken by friends, etc. For labels, Zeng [31] scan over the 81 tags of each image and manually merged them into 7 classes. Each image belongs to one of the 7 classes. Reddit [8] is a graph dataset constructed from Reddit posts. In this case, the node label is the community or subreddit that a post belongs to.

These datasets involve both the transductive task and inductive task. The transductive task in our experiments is semi-supervised node classification on one graph; the inductive task is the node classification on multiple graphs. The information of them is presented in Table I. The label rate in the table means the ratio of training data.

2) GNNs Models for Training: We have applied the RWT to train GCN and GAT models, respectively, which both are representative and widely used GNN models. The default models contain two layers. The hidden layers involve different sizes based on different datasets: for Cora, Citeseer, and Pubmed, the hidden layers' size is 32; for Flickr and Reddit, the hidden size is 128 for GCN layer and 8 for GAT layer. The dropout rate is 0.5 for Cora, Citeseer, Pubmed, and 0.1 for Flickr, Reddit. We employ the ADAM [11] as the optimizer. The learning rate is 0.01, with weight decay as 5×10^{-4}.

3) Comparison Methods: We compare the GNNs trained by RWT with state-of-the-art baseline methods:

Comparison Models

- **GCN** [12]: GCN is a semi-supervised method for the node classification, which operates on the whole graph.
- **GAT** [27]: GAT is an attention-based graph neural network for the node classification.
- **GraphSAGE** [8]: GraphSAGE is a general inductive framework that leverages node feature information to generate node embeddings for unseen data efficiently.
- **Cluster-GCN** [5]: Cluster-GCN is a suitable framework for SGD-based training. It samples a block of nodes that associate with a dense subgraph identified by a graph clustering algorithm and restricts the neighborhood search within this subgraph.

Comparison Samplers

TABLE II
TEST ACCURACY RESULTS ON ALL DATASETS

Methods	Transductive			Inductive	
	Cora	Citeseer	Pubmed	Flickr	Reddit
GraphSAGE	0.7660	0.6750	0.7610	0.4030	0.9300
Cluster-GCN	0.682	0.628	0.7947	0.4097	**0.9523**
GCN	0.815	0.7030	0.7890	0.4400	0.9333
GCN + Random	0.7945	0.687	0.7345	0.4713	0.8243
GCN + BFS	0.8144	0.7079	0.7971	0.4754	0.8123
GCN + RWT	**0.825**	**0.7127**	**0.8259**	**0.4797**	0.9495
GAT	**0.8300**	0.7130	0.7903	-	-
GAT + Random	0.7921	0.6607	0.6765	0.4534	0.6452
GAT + BFS	0.7756	0.6500	0.7080	0.4642	0.7297
GAT + RWT	0.7994	**0.7212**	**0.8210**	**0.4724**	**0.8699**

"-" insufficient memory.

- **Ripple Walk Sampler**: Ripple Walk Sampler is the sampler proposed in this paper.
- **Random**: The Random sampler randomly selects a certain number from all nodes to form a subgraph for training.
- **BFS**: The BFS sampler performs a breadth-first search from the starting node to select subgraphs.

For the Cluster-GCN, the number of clusters is set to make the average subgraph size the same as Ripple Walk Sampler. Meanwhile, the self-comparison is conducted among Ripple Walk Sampler, BFS sampling, and Random sampling strategies. For these samplers, the sampled subgraph size on Cora and Citeseer datasets is $S = 1500$, on Pubmed and Flickr $S = 3000$, on Reddit $S = 5000$. For the subgraph mini-batch size M, it can be varying for different datasets. Generally, we make the overall sampled nodes at least twice of the nodes in a full graph, which also means $M \times S \simeq 10|\mathcal{V}|$. In the following parts, the "GCN + Random / BFS / RWT" denotes GCN model training with subgraphs from Random sampler, BFS sampler and Ripple Walk Sampler, respectively.

We run the experiments on the Server with 3 GTX-1080 ti GPUs, and all codes are implemented in Python. Code is available at: https://github.com/anonymous2review/RippleWalk.

B. Experimental Results with Analysis

1) Task Performance Analysis: The overall task performance of RWT and comparison methods are exhibited in Table IV-B. The most important thing about the training framework RWT is to ensure the training effect of GNN. On the premise of competitiveness or better task performance, the advantages of time and space can be meaningful. We can first observe that GCN and GAT with RWT outperforms plain GCN and GAT in most of the cases. Here, both GCN and GAT contain two layers. For the GCN model, RWT has better overall performance than GraphSAGE and Cluster-GCN; for the GAT, even in some cases when training with a full graph cannot be executed due to limited memory space (e.g., on Flickr and Reddit), GAT with RWT can successfully run and achieve high performance. For the self-comparison, Ripple Walk Sampler achieves the best results compared with random and BFS sampling. Generally speaking, GNNs with RWT can achieve the same level or even better testing performance compared with other popular baseline methods. Meanwhile, the advantages of RWT on space complexity are significant. In the following

TABLE III
MEMORY SPACE USAGE (THE UNIT IS MB)

	Cora	Citeseer	Pubmed	Flickr	Reddit
GCN	535	605	2057	30392	212003
GCN + RWT	**509**	**587**	**1235**	**922**	**1101**
GAT	6921	10277	11868	243089	243089
GAT + RWT	**2121**	**2469**	**2629**	**12000**	**12080**

part, we will answer four evaluation questions mentioned before to validate the efficiency of RWT.

2) Space-consuming Analysis: One of the critical advantages of RWT compared with plain GNNs training is less space-consuming. To answer **Question 1**, we first compare the memory space usage and show them in Table III. It is obvious that training GNNs with RWT requires less memory space than plain GNNs. Especially for GAT, the space usage of RWT is much less than using the full graph. Therefore, when training plain GNNs is too space-consuming to run, RWT can help conduct the training process of GNNs and the performance can be guaranteed. Meanwhile, the less space-consuming of RWT enables the GNNs to be carried out on GPUs, which can further accelerate the training. In this part, we do not compare RWT with other baseline methods, since for the neighbor node sampling methods such as GraphSAGE [8], FastGCN [4] or methods in [3], [10] do not apply the subgraph training concept. Thus their space complexity will be the same with the plain GNNs. Therefore we can conclude that while maintaining the high-level testing performance, RWT can significantly minimize the space-consuming when running GNNs. In this way, RWT break the memory bottleneck when facing huge graphs. The experimental results also validate our previous theoretic analyses about both effectiveness and efficiency.

TABLE IV
TRAINING TIME (THE UNIT IS SECOND)

	Cora	Citeseer	Pubmed	Flickr	Reddit
GCN	4.573	1.968	61.90	1161.92	25370
GCN + RWT	**1.964**	**1.826**	**8.698**	**1.179**	**7.722**
GAT	413.3	500.1	-	-	-
GAT + RWT	**71.44**	**47.06**	**139.4**	**68.06**	**2614**

"-" insufficient memory.

3) Time-consuming Analysis: To answer **Question 2**, we further record the duration time of the training process and present it in Table IV. The running time of GCN on Flickr and Reddit datasets, GAT + RWT on Reddit dataset is based on the CPU server. All other convergence time is recorded on GPUs. By applying RWT, the running time of GNNs models can be reduced by a large margin. Thus the RWT on GNNs not only requires less memory space but also can accelerate the convergence of the training process. Such results can be explained intuitively since each training epoch of RWT involves less space and computation complexity. While sharing the same computational quantity, RWT possesses more training iterations than plain GNNs training.

4) Sampler performance Analysis: To answer **Question 3**, we conduct comparison experiments among different samplers. From Table IV-B, we can notice that Ripple Walk Sampler outperforms other samplers on all five datasets. These results verify that the subgraphs provided by Ripple Walk Sampler are more beneficial for graph model training. Random sampler

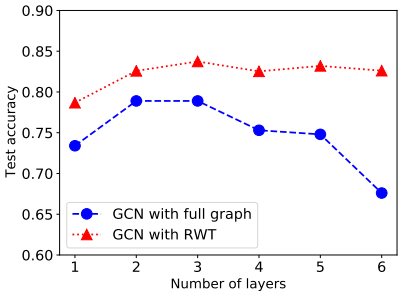

Fig. 2. GCN with Deeper Layers

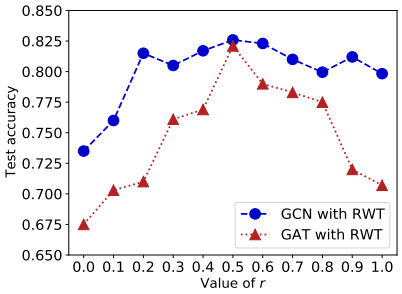

Fig. 3. GNNs with Different r

leans towards randomness, while BFS sampler focuses on connectivity too much. They can not consider randomness and connectivity simultaneously, yet Ripple Walk Sampler can do so. It also validates our theoretical analysis in Section III-B.

5) For Deeper Graph Networks: We test GCN models with different numbers of layers to answer **Question 4**. The results are shown in Figure 2. We show the results on the Pubmed dataset, and the experimental results are consistent in other datasets. We can observe that GCN with RWT achieves better performance than plain GCN on the test loss and accuracy. Besides, with the structure goes deeper, even when the performance of GCN decreases, GCN with RWT achieves higher performance. Thus with the support of RWT, the problem of *oversmoothing* can be eliminated, and GNNs models can be designed with deeper structure. From such phenomenon, deeper GNNs models such as [32] can benefit from the RWT.

C. Parameter Analysis

1) Expansion Ratio r Analysis: To verify the analysis of r in Subsection III-C, we implement experiments of Ripple Walk Sampler with different expansion ratios, and present the results in Figure 3. We show the results on Pubmed, and the experimental results are consistent in all datasets. According to previous analysis in Subsection III-C, $r \to 0$ or $r \to 1$ do not help maintain the randomness and connectivity characteristics in subgraphs. From the results we can observe that when $r = 0.5$, RWT achieve the best performance and the performance decreases when $r \to 1$ or $r \to 0$. Therefore, the results verify our previous analysis, The subgraphs sampled by Ripple Walk Sampler consider both randomness and connectivity, which are beneficial to subgraph-based training for GNNs.

2) Subgraph Mini-batch Size M Analysis: As mentioned above, we have applied the subgraph mini-batch size to satisfy the $M \times S \simeq 10|\mathcal{V}|$. Give a fixed subgraph size S, we

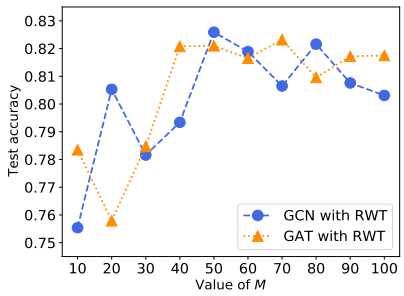

Fig. 4. GNNs with Different values of M

also attempt different values of M and present the results on Pubmed dataset with $S = 3000$ and $M \in \{10, 20, \ldots, 100\}$ in Figure 4. Both GCN and GAT with RWT have relatively worse performance when M is small (e.g., $M \in \{10, 20, 30\}$). With the increase of M, the performance becomes better and finally maintains at a high level. It is intuitive that with a smaller subgraph mini-batch size M, less information of the full graph can be covered. While M rising, more subgraphs combinations will be selected to draw the information (e.g., connections) within the full graph. On the other hand, along with increasing M, subgraphs with similar nodes and edges could be sampled. Such similar subgraphs may cause redundant information extraction from the full graph, and the performance will finally converge instead of improving continuously.

V. Conclusion

In this paper, we have introduced a subgraph-based training framework RWT for GNNs, which combines the idea of training GNNs with mini-batch subgraphs and a novel subgraph sampling method Ripple Walk Sampler. We analyze the effectiveness and efficiency of the Ripple Walk and prove them from the theoretical perspective. Extensive experiments demonstrate that RWT achieves the same level or even better performance, but less training time and device memory space are required. At the same time, RWT can help relieve the problem of *oversmoothing* when models go deeper, enabling the GNNs to have stronger learning power and potential.

References

[1] George W. Brown. On small-sample estimation. *The Annals of Mathematical Statistics*, 18(4):582–585, 1947.

[2] Joan Bruna, Wojciech Zaremba, Arthur Szlam, and Yann LeCun. Spectral networks and locally connected networks on graphs. *arXiv:1312.6203*, 2013.

[3] Jianfei Chen, Jun Zhu, and Le Song. Stochastic training of graph convolutional networks with variance reduction. In *Proceedings of the 35th International Conference on Machine Learning*, 2018.

[4] Jie Chen, Tengfei Ma, and Cao Xiao. Fastgcn: fast learning with graph convolutional networks via importance sampling. *arXiv:1801.10247*, 2018.

[5] Wei-Lin Chiang, Xuanqing Liu, Si Si, Yang Li, Samy Bengio, and Cho-Jui Hsieh. Cluster-gcn: An efficient algorithm for training deep and large graph convolutional networks. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 257–266, 2019.

[6] Alex Fout, Jonathon Byrd, Basir Shariat, and Asa Ben-Hur. Protein interface prediction using graph convolutional networks. In *Advances in neural information processing systems*, pages 6530–6539, 2017.

[7] Hongyang Gao, Zhengyang Wang, and Shuiwang Ji. Large-scale learnable graph convolutional networks. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1416–1424, 2018.

[8] Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. In *Advances in neural information processing systems*, pages 1024–1034, 2017.

[9] Mikael Henaff, Joan Bruna, and Yann LeCun. Deep convolutional networks on graph-structured data. *arXiv:1506.05163*, 2015.

[10] Wenbing Huang, Tong Zhang, Yu Rong, and Junzhou Huang. Adaptive sampling towards fast graph representation learning. In *Advances in neural information processing systems*, pages 4558–4567, 2018.

[11] Diederik P. Kingma and Jimmy Lei Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representaion*, 2015.

[12] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representaion*, 2017.

[13] Ron Levie, Federico Monti, Xavier Bresson, and Michael M Bronstein. Cayleynets: Graph convolutional neural networks with complex rational spectral filters. *IEEE Transactions on Signal Processing*, 67(1):97–109, 2018.

[14] Qimai Li, Zhichao Han, and Xiao-Ming Wu. Deeper insights into graph convolutional networks for semi-supervised learning. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.

[15] Ruoyu Li, Sheng Wang, Feiyun Zhu, and Junzhou Huang. Adaptive graph convolutional neural networks. In *Thirty-second AAAI conference on artificial intelligence*, 2018.

[16] Renjie Liao, Zhizhen Zhao, Raquel Urtasun, and Richard S Zemel. Lanczosnet: Multi-scale deep graph convolutional networks. *arXiv:1901.01484*, 2019.

[17] László Lovász et al. Random walks on graphs: A survey. *Combinatorics, Paul erdos is eighty*, 2(1):1–46, 1993.

[18] Julian McAuley and Jure Leskovec. Image labeling on a network: using social-network metadata for image classification. In *European conference on computer vision*, pages 828–841. Springer, 2012.

[19] Alessio Micheli. Neural network for graphs: A contextual constructive approach. *IEEE Transactions on Neural Networks*, 20(3):498–511, 2009.

[20] Federico Monti, Davide Boscaini, Jonathan Masci, Emanuele Rodola, Jan Svoboda, and Michael M Bronstein. Geometric deep learning on graphs and manifolds using mixture model cnns. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5115–5124, 2017.

[21] Vinod Nair and Geoffrey E Hinton. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th international conference on machine learning (ICML-10)*, pages 807–814, 2010.

[22] Yuxiang Ren and Jiawei Zhang. Hgat: Hierarchical graph attention network for fake news detection. *arXiv preprint arXiv:2002.04397*, 2020.

[23] Yu Rong, Wenbing Huang, Tingyang Xu, and Junzhou Huang. Dropedge: Towards deep graph convolutional networks on node classification. *arXiv:1907.10903*, 2019.

[24] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1):61–80, 2008.

[25] Prithviraj Sen, Galileo Namata, Mustafa Bilgic, Lise Getoor, Brian Galligher, and Tina Eliassi-Rad. Collective classification in network data. *AI magazine*, 29(3):93–93, 2008.

[26] Alessandro Sperduti and Antonina Starita. Supervised neural networks for the classification of structures. *IEEE Transactions on Neural Networks*, 8(3):714–735, 1997.

[27] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks. In *International Conference on Learning Representaion*, 2018.

[28] Quan Wang, Zhendong Mao, Bin Wang, and Li Guo. Knowledge graph embedding: A survey of approaches and applications. *IEEE Transactions on Knowledge and Data Engineering*, 29(12):2724–2743, 2017.

[29] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? *arXiv:1810.00826*, 2018.

[30] Keyulu Xu, Chengtao Li, Yonglong Tian, Tomohiro Sonobe, Ken-ichi Kawarabayashi, and Stefanie Jegelka. Representation learning on graphs with jumping knowledge networks. In *Proceedings of the 35th International Conference on Machine Learning*, 2018.

[31] Hanqing Zeng, Hongkuan Zhou, Ajitesh Srivastava, Rajgopal Kannan, and Viktor Prasanna. Graphsaint: Graph sampling based inductive learning method. *arXiv preprint arXiv:1907.04931*, 2019.

[32] Jiawei Zhang and Lin Meng. Gresnet: Graph residual network for reviving deep gnns from suspended animation. In *arXiv:1909.05729*, 2019.

[33] Lingxiao Zhao and Leman Akoglu. Pairnorm: Tackling oversmoothing in gnns. *arXiv:1909.12223*, 2019.