

---

# LazyGNN: Large-Scale Graph Neural Networks via Lazy Propagation

---

Rui Xue<sup>1</sup> Haoyu Han<sup>2</sup> MohamadAli Torkamani<sup>3</sup> Jian Pei<sup>4</sup> Xiaorui Liu<sup>1</sup>

## Abstract

Recent works have demonstrated the benefits of capturing long-distance dependency in graphs by deeper graph neural networks (GNNs). But deeper GNNs suffer from the long-lasting scalability challenge due to the neighborhood explosion problem in large-scale graphs. In this work, we propose to capture long-distance dependency in graphs by shallower models instead of deeper models, which leads to a much more efficient model, LazyGNN, for graph representation learning. Moreover, we demonstrate that LazyGNN is compatible with existing scalable approaches (such as sampling methods) for further accelerations through the development of mini-batch LazyGNN. Comprehensive experiments demonstrate its superior prediction performance and scalability on large-scale benchmarks. The implementation of LazyGNN is available at [https://github.com/RXPHD/Lazy\\_GNN](https://github.com/RXPHD/Lazy_GNN).

## 1. Introduction

Graph neural networks (GNNs) have been widely used for representation learning on graph-structured data (Hamilton, 2020; Ma & Tang, 2021), and they achieve promising state-of-the-art performance on various general graph learning tasks, such as node classification, link prediction, and graph classification (Kipf & Welling, 2016; Gasteiger et al., 2019; Veličković et al., 2017; Wu et al., 2019) as well as a variety of important applications, such as recommendation systems, social network analysis, and transportation prediction. In particular, recent research in deeper GNNs has generally revealed the performance gains from capturing long-distance relations in graphs by stacking more graph convolution layers or unrolling various fixed point iterations (Gasteiger et al., 2018; Gu et al., 2020; Liu et al.,

2020; Chen et al., 2020a; Li et al., 2021; Ma et al., 2020; Pan et al., 2020; Zhu et al., 2021; Chen et al., 2020b). However, the recursive feature propagations in deeper GNNs lead to the well-known neighborhood explosion problem since the number of neighbors grows exponentially with the number of feature propagation layers (Hamilton et al., 2017; Chen et al., 2018a). This causes tremendous scalability challenges for data sampling, computation, memory, parallelism, and end-to-end training when employing GNNs on large-scale graphs. It greatly limits GNNs’ broad applications in large-scale industry-level applications due to limited computation and memory resources (Ying et al., 2018; Shao et al., 2022).

A large body of existing research improves the scalability and efficiency of large-scale GNNs using various innovative designs, such as sampling methods, pre-computing or post-computing methods, and distributed methods. Although these approaches mitigate the neighborhood explosion problem, they still face various limitations when they are applied to deeper GNNs. For instance, sampling approaches (Hamilton et al., 2017; Chen et al., 2018a; Zeng et al., 2020; Zou et al., 2019; Fey et al., 2021; Yu et al., 2022) usually incur large approximation errors and suffer from performance degradation as demonstrated in large-scale OGB benchmarks or require large additional memory or storage space to store activation values in hidden layers. Pre-computing or post-computing methods (Wu et al., 2019; Rossi et al., 2020; Sun et al., 2021; Zhang et al., 2022; Bojchevski et al., 2020; Zhu, 2005; Huang et al., 2020) lose the benefits of end-to-end training and usually suffer from performance loss. Distributed methods (Chiang et al., 2019; Chai et al., 2022; Shao et al., 2022) distribute large graphs to multiple servers for parallel training, but they either neglect the inter-server edges or suffer from expensive feature communication between servers.

In this work, we take a substantially different and novel perspective and propose to capture long-distance dependency in graphs by shallower GNNs instead of deeper ones. The key intuition comes from the fact that node information will be propagated over the graph many times during the training process so we may only need to propagate information lazily by reusing the propagated information over the training iterations. This intuition leads to the proposed LazyGNN that solves the inherent neighborhood explosion problem by significantly reducing the number of aggregation layers while

---

<sup>1</sup>North Carolina State University, Raleigh, US <sup>2</sup>Michigan State University, East Lansing, US <sup>3</sup>Amazon, US (this work does not relate to the author’s position at Amazon) <sup>4</sup>Duke University, Durham, US. Correspondence to: Xiaorui Liu <xliu96@ncsu.edu>.

still capturing long-distance dependency in graphs through lazy propagation. Multiple technical challenges such as the risk of over-smoothing, additional variation due to feature dropout, and back-propagation through historical computation graphs are addressed through innovative designs in LazyGNN. Moreover, since LazyGNN is a shallow model, it naturally tackles the limitations that existing scalable approaches suffer from when handling large-scale and deeper GNNs. Therefore, the contribution of LazyGNN is orthogonal and complementary to existing efforts, and various scalable approaches can be used in LazyGNN for further acceleration. We demonstrate this by developing a highly scalable and efficient mini-batch LazyGNN based on sampling methods. To the best of our knowledge, LazyGNN is the first GNN model being versatilely friendly to data sampling, computation, memory, parallelism, and end-to-end training but still captures long-distance dependency in graphs. Our contributions can be summarized as follows:

- We reveal a novel perspective to solve the neighborhood explosion problem by exploiting the computation redundancy in training GNNs;
- A novel shallow model, LazyGNN, is proposed to capture long-distance dependency in graphs for end-to-end graph representation learning through lazy forward and backward propagations;
- We demonstrate that existing scalable approaches can be compatible with LazyGNN by introducing a highly efficient and scalable mini-batch LazyGNN to handle large-scale graphs based on sampling methods.
- Comprehensive experiments and studies demonstrate that LazyGNN achieves superior prediction performance and efficiency on large-scale graph datasets.

## 2. Preliminary

**Notations.** A graph is represented by  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  where  $\mathcal{V} = \{v_1, \dots, v_N\}$  is the set of nodes and  $\mathcal{E} = \{e_1, \dots, e_M\}$  is the set of edges. Suppose that each node is associated with a  $d$ -dimensional feature vector, and the original features for all nodes are denoted as  $\mathbf{X}_{\text{fea}} \in \mathbb{R}^{N \times d}$ . The graph structure of  $\mathcal{G}$  can be represented by an adjacency matrix  $\mathbf{A} \in \mathbb{R}^{N \times N}$ , where  $\mathbf{A}_{ij} > 0$  when there exists an edge between node  $v_i$  and  $v_j$ , otherwise  $\mathbf{A}_{i,j} = 0$ . The symmetrically normalized graph Laplacian matrix is defined as  $\tilde{\mathbf{L}} = \mathbf{I} - \tilde{\mathbf{A}}$  with  $\tilde{\mathbf{A}} = \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}$  where  $\mathbf{D}$  is the degree matrix. Next, we briefly introduce the graph signal denoising and fixed point iteration perspective of GNNs that provide a better understanding of the computation trajectory and long-distance dependency in graphs. Finally, we provide a preliminary study to reveal the computation redundancy in GNNs that motivates the development of LazyGNN.

### 2.1. GNNs as Graph Signal Denoising

It is recently established that many popular GNNs can be uniformly understood as solving graph signal denoising problems with various diffusion properties and that the long-distance dependency can be well captured by unrolling various fixed point iterations (Ma et al., 2020; Pan et al., 2020; Zhu et al., 2021; Chen et al., 2020b; Gu et al., 2020). For instance, the message passing in GCN (Kipf & Welling, 2016),  $\mathbf{X}_{\text{out}} = \tilde{\mathbf{A}} \mathbf{X}_{\text{in}}$ , can be considered as one gradient descent iteration for minimizing  $\text{tr}(\mathbf{X}^\top (\mathbf{I} - \tilde{\mathbf{A}}) \mathbf{X})$  with the initialization  $\mathbf{X}_0 = \mathbf{X}_{\text{in}}$ . The message passing scheme in APPNP (Klicpera et al., 2018) follows the aggregation  $\mathbf{X}_{l+1} = (1 - \alpha) \tilde{\mathbf{A}} \mathbf{X}_l + \alpha \mathbf{X}_{\text{in}}$  that iteratively minimizes  $\|\mathbf{X} - \mathbf{X}_{\text{in}}\|_F^2 + (1/\alpha - 1) \text{tr}(\mathbf{X}^\top (\mathbf{I} - \tilde{\mathbf{A}}) \mathbf{X})$  with the initialization  $\mathbf{X}_0 = \mathbf{X}_{\text{in}}$  where  $l$  is the index of layers. Implicit GNN (Gu et al., 2020) adopts projected gradient descent to solve the fixed point problem  $\mathbf{X} = \phi(\mathbf{W} \mathbf{X} \tilde{\mathbf{A}} + \mathbf{B})$ . Please refer to the reference (Ma et al., 2020; Pan et al., 2020; Zhu et al., 2021; Chen et al., 2020b) for such understanding of many other popular GNN models. Moreover, a large number of advanced GNN models have been inspired from this perspective (Chen et al., 2022; Liu et al., 2021a;b; Yang et al., 2021a;b; Jia & Benson, 2022; Jiang et al., 2022).

### 2.2. Computation Redundancy in Training GNNs

In the training process of GNNs, the model parameters are updated following gradient descent style algorithms such as Adam (Kingma & Ba, 2014). Therefore, the model as well as hidden features in GNNs changes smoothly, especially in the late stages when both the learning rate and gradient norm diminish. This intuition motivates us to investigate the computation redundancy in GNNs. Specifically, we measure the relative changes of the hidden features in the last layer ( $L$ -th layer) between epochs  $k + 1$  and  $k$ , i.e.,  $\|\mathbf{X}_L^{k+1} - \mathbf{X}_L^k\|_F / \|\mathbf{X}_L^k\|_F$ , over the training iterations on Cora and Pubmed dataset (Kipf & Welling, 2016) using representative models such as GCN (Kipf & Welling, 2016) and APPNP (Klicpera et al., 2018). We show two cases in Figure 1 and Figure 2, depending on whether dropout layers are used.

The relative changes of hidden features shown in Figure 1 and Figure 2 demonstrate the following: (1) when there is no dropout, the hidden features barely change as the training goes; (2) if dropout is used, it will incur additional variation in the hidden features due to the randomness in dropout layers. Both cases demonstrate that the activation values in hidden layers of GNNs indeed change slowly, indicating the existence of *computation redundancy*: the computation in successive training iterations is highly similar. This observation not only validates the rationality of historical embedding used in existing works such as VR-GNN (Chen et al., 2017) and GNNAutoScale (Fey et al., 2021) but also

motivates the lazy propagation in this work.




Figure 1. Feature changes  $\|\mathbf{X}^{k+1} - \mathbf{X}^k\|_F / \|\mathbf{X}^k\|_F$  on Cora.




Figure 2. Feature changes  $\|\mathbf{X}^{k+1} - \mathbf{X}^k\|_F / \|\mathbf{X}^k\|_F$  on Pubmed.

### 3. Lazy Graph Neural Networks

In this section, we propose a novel shallow LazyGNN that uses a few aggregation layers to capture long-distance dependency in graphs through lazy forward and backward propagations. Then a mini-batch LazyGNN is introduced to handle large-scale graphs with efficient data sampling, computation, and memory usage. Complexity analyses are also provided to illustrate the superior scalability of LazyGNN.

#### 3.1. GNNs with Lazy Propagation

Existing research has demonstrated the benefits of capturing long-distance relations in graphs by stacking more feature aggregation layers or unrolling various fixed point iterations as introduced in Section 1. However, these deeper GNNs are not efficient due to the neighborhood explosion problem. Our preliminary study in Section 2 reveals the computation redundancy in GNNs over the training iterations: the hidden features in GNNs evolve slowly such that the computation of feature aggregations is highly correlated and redundant over training iterations. This observation motivates us to develop a novel shallow LazyGNN that captures long-distance relations in graphs by propagating information lazily with very few message-passing layers.

Without loss of generality, we illustrate the main idea of lazy propagations using the most common and widely used graph signal denoising problem (Zhou et al., 2003; Ma et al., 2020; Yang et al., 2021a):

$$\min_{\mathbf{X}} \|\mathbf{X} - \mathbf{X}_{\text{in}}\|_F^2 + (1/\alpha - 1) \text{tr}(\mathbf{X}^\top (\mathbf{I} - \tilde{\mathbf{A}})\mathbf{X}), \quad (1)$$

where the first term maintains the proximity with node hidden features  $\mathbf{X}_{\text{in}}$  after feature transformation, and the second Laplacian smoothing regularization encodes smoothness assumption on graph representations. Note that we only take this as an example to illustrate the main idea of lazy propagation, but the idea can be applied to other GNNs inspired by any denoising problems or fixed point iterations with different properties and design motivations. Next, we illustrate the lazy forward and backward propagations in LazyGNN as well as the innovative designs that solve the technical challenges.

**Forward Computation.** From Eq. (1), we can derive the high-order iterative graph diffusion as in APPNP (Klicpera et al., 2018) with  $f$  being the feature transformation:

$$\mathbf{X}_{\text{in}}^k = f(\mathbf{X}_{\text{fea}}, \Theta^k), \quad (2)$$

$$\mathbf{X}_0^k = \mathbf{X}_{\text{in}}^k, \quad (3)$$

$$\mathbf{X}_{l+1}^k = (1 - \alpha)\tilde{\mathbf{A}}\mathbf{X}_l^k + \alpha\mathbf{X}_{\text{in}}^k, \quad \forall l = 0, \dots, L - 1 \quad (4)$$

where  $l$  and  $k$  denote the index of layers and training iterations, respectively. The key insight of LazyGNN is that the approximate solution of Eq. (1) (i.e.,  $\mathbf{X}_L^k$ ) evolves smoothly since  $\mathbf{X}_{\text{in}}^k = f(\mathbf{X}_{\text{fea}}, \Theta^k)$  changes smoothly with model parameters  $\Theta^k$ . Intuitively, the features have been propagated over the graph many times in previous training iterations so it suffices to propagate features lazily by reusing existing computation.

Formally, we propose LazyGNN to leverage the computation redundancy between training iterations by mixing the diffusion output in iteration  $k - 1$  (i.e.,  $\mathbf{X}_L^{k-1}$ ) into the initial embedding of the diffusion process in training iteration  $k$ , namely  $\mathbf{X}_0^k = (1 - \beta)\mathbf{X}_L^{k-1} + \beta\mathbf{X}_{\text{in}}^k$ . This is because, in practice, dropout is commonly used in deep learning to prevent overfitting, which introduces additional variations in the feature embedding as demonstrated in Figure 1 and Figure 2. Therefore, we introduce this momentum correction to compensate for such disturbance by mixing current and history features with hyperparameter  $\beta$ . In practice, small  $\beta$  is favored if the dropout rate is small. To summarize, the forward computation of LazyGNN works as follows:

$$\mathbf{X}_{\text{in}}^k = f(\mathbf{X}_{\text{fea}}, \Theta^k), \quad (5)$$

$$\mathbf{X}_0^k = (1 - \beta)\mathbf{X}_L^{k-1} + \beta\mathbf{X}_{\text{in}}^k, \quad (6)$$

$$\mathbf{X}_{l+1}^k = (1 - \alpha)\tilde{\mathbf{A}}\mathbf{X}_l^k + \alpha\mathbf{X}_{\text{in}}^k, \quad \forall l = 0, \dots, L - 1 \quad (7)$$

By lingering the computation over training iterations as shown in Figure 3, the feature aggregation layers in Eq. (7) solve the denoising problem in Eq. (1) with an implicitly

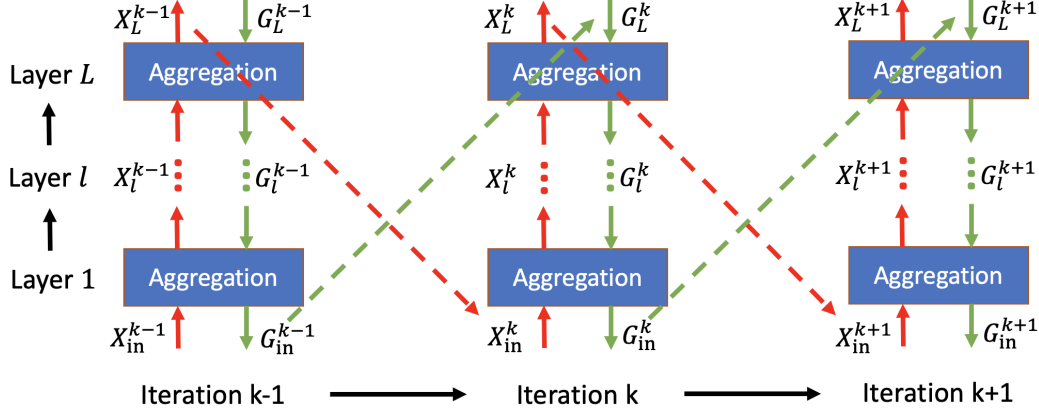


Figure 3. LazyGNN with Lazy **forward** (red) and **backward** (green) propagations.

large number of steps although there are only  $L$  layers in each training iteration. Therefore, it only requires a few feature aggregation layers (a small  $L$ ) to approximate the fixed point solution  $\mathbf{X}_*^k$  of Eq. (1). Note that we find  $L = 1$  and  $L = 2$  work very well in our experiments in Section 4. In the extreme case when the learning rate and dropout rate are 0, the forward computation of LazyGNN is equivalent to running forward propagations  $L \times K$  times continuously with  $K$  being the total number of training iterations.

**Remark 1 (Long-distance dependency).** Through lazy propagation, LazyGNN is able to capture long-distance dependency in graphs with a small number of feature aggregation layers because LazyGNN accumulatively propagates information to distant nodes in the graphs over many training iterations. In contrast, existing works try to capture long-distance dependency in graphs by increasing the number of feature propagation layers, which is less efficient.

**Remark 2 (Over-smoothing).** In contrast to many GNN models such as GCN or GAT that suffer from the over-smoothing issue when more layers are stacked, the accumulation of feature aggregations over training iterations in LazyGNN will not cause the over-smoothing issue because the residual connection  $\mathbf{X}_{in}^k$  in Eq. (7) determines the fixed point and prevents the over-smoothed solution, as will be verified in Section 4.3.

**Backward Computation.** While it is feasible to leverage the computation from previous training iterations in the forward computation, it is highly non-trivial to compute the gradient for the model update in the backward computation since the computation graphs from previous training iterations have been destroyed and released in the memory. In other words, there is no way to directly compute the backpropagation through the history variables  $\{\mathbf{X}_L^{k-1}, \mathbf{X}_L^{k-2}, \dots\}$  in current iterations  $k$  as being done in sequential models such as RNN (Hochreiter & Schmidhuber, 1997) and Transformer (Vaswani et al., 2017). Therefore,

we choose to compute the gradient indirectly via the implicit function theorem (Bai et al., 2019; Gu et al., 2020).

**Theorem 1 (Implicit Gradient).** Let  $\mathbf{X}_*$  be the fixed point solution of function  $g(\mathbf{X}_*, \mathbf{X}_{in})$ , i.e.,  $g(\mathbf{X}_*, \mathbf{X}_{in}) = \mathbf{0}$ . Given the gradient of loss function  $\mathcal{L}(\mathbf{X}_*, \mathbf{Y})$  with respect to the fixed point  $\mathbf{X}_*$ , i.e.,  $\frac{\partial \mathcal{L}}{\partial \mathbf{X}_*}$ , the gradient of loss  $\mathcal{L}$  with respect to feature  $\mathbf{X}_{in}$  can be computed as:

$$\frac{\partial \mathcal{L}}{\partial \mathbf{X}_{in}} = -\frac{\partial \mathcal{L}}{\partial \mathbf{X}_*} (\mathbf{J}|_{\mathbf{X}_*})^{-1} \frac{\partial g(\mathbf{X}_*, \mathbf{X}_{in})}{\partial \mathbf{X}_{in}} \quad (8)$$

where  $\mathbf{J}|_{\mathbf{X}_*} = \frac{\partial g(\mathbf{X}_*, \mathbf{X}_{in})}{\partial \mathbf{X}_*}$  is the Jacobian matrix of  $g(\mathbf{X}_*, \mathbf{X}_{in})$  evaluated at  $\mathbf{X}_*$ .

*Proof.* Take the first-order derivative on both sides of the fixed point equation  $g(\mathbf{X}_*, \mathbf{X}_{in}) = \mathbf{0}$ :

$$\frac{\partial g(\mathbf{X}_*, \mathbf{X}_{in})}{\partial \mathbf{X}_{in}} + \frac{\partial g(\mathbf{X}_*, \mathbf{X}_{in})}{\partial \mathbf{X}_*} \frac{d\mathbf{X}_*}{d\mathbf{X}_{in}} = \mathbf{0}$$

$$\frac{d\mathbf{X}_*}{d\mathbf{X}_{in}} = -\left(\frac{\partial g(\mathbf{X}_*, \mathbf{X}_{in})}{\partial \mathbf{X}_*}\right)^{-1} \frac{\partial g(\mathbf{X}_*, \mathbf{X}_{in})}{\partial \mathbf{X}_{in}}$$

Using  $\frac{\partial \mathcal{L}}{\partial \mathbf{X}_{in}} = \frac{\partial \mathcal{L}}{\partial \mathbf{X}_*} \frac{d\mathbf{X}_*}{d\mathbf{X}_{in}}$ , we obtain:

$$\frac{\partial \mathcal{L}}{\partial \mathbf{X}_{in}} = -\frac{\partial \mathcal{L}}{\partial \mathbf{X}_*} \left(\frac{\partial g(\mathbf{X}_*, \mathbf{X}_{in})}{\partial \mathbf{X}_*}\right)^{-1} \frac{\partial g(\mathbf{X}_*, \mathbf{X}_{in})}{\partial \mathbf{X}_{in}}.$$

□

Specifically, for the problem in Eq. (1), we have the fixed point equation:

$$g(\mathbf{X}_*, \mathbf{X}_{in}) = \mathbf{X}_* - \mathbf{X}_{in} + \left(\frac{1}{\alpha} - 1\right)(\mathbf{I} - \tilde{\mathbf{A}})\mathbf{X}_* = \mathbf{0} \quad (9)$$

which gives  $\frac{\partial g(\mathbf{X}_*, \mathbf{X}_{in})}{\partial \mathbf{X}_{in}} = -\mathbf{I}$  and  $\mathbf{J}|_{\mathbf{X}_*} = \frac{1}{\alpha}(\mathbf{I} - (1 - \alpha)\tilde{\mathbf{A}})$ .

Therefore, according to Theorem 1, the gradient of loss  $\mathcal{L}$  with respect to  $\mathbf{X}_{\text{in}}$  can be computed as follows:

$$\frac{\partial \mathcal{L}}{\partial \mathbf{X}_{\text{in}}} = \alpha \frac{\partial \mathcal{L}}{\partial \mathbf{X}_*} \left( \mathbf{I} - (1 - \alpha) \tilde{\mathbf{A}} \right)^{-1}. \quad (10)$$

However, it is still infeasible to compute the expensive matrix inversion so we propose to approximate it by the iterative backward gradient propagation:

$$\mathbf{G}_L = \frac{\partial \mathcal{L}}{\partial \mathbf{X}_L} \left( \approx \frac{\partial \mathcal{L}}{\partial \mathbf{X}_*} \right) \quad (11)$$

$$\mathbf{G}_l = (1 - \alpha) \tilde{\mathbf{A}} \mathbf{G}_{l+1} + \alpha \frac{\partial \mathcal{L}}{\partial \mathbf{X}_L} \quad \forall l = L - 1, \dots, 0 \quad (12)$$

where  $\frac{\partial \mathcal{L}}{\partial \mathbf{X}_*}$  is approximated by  $\frac{\partial \mathcal{L}}{\partial \mathbf{X}_L}$ , and  $\mathbf{G}_0$  provides an approximation for gradient  $\frac{\partial \mathcal{L}}{\partial \mathbf{X}_{\text{in}}}$ . It is clear that the backward computation requires gradient propagation over the graph, which is symmetric to and as expensive as the vanilla forward computation. Similarly, to reduce the number of gradient propagation layers, we propose to propagate the gradient lazily by leveraging the gradient  $\frac{\partial \mathcal{L}}{\partial \mathbf{X}_{\text{in}}^{k-1}}$  computed in the previous training iteration  $k - 1$  as shown in Figure 3:

$$\mathbf{G}_L^k = (1 - \gamma) \frac{\partial \mathcal{L}}{\partial \mathbf{X}_{\text{in}}^{k-1}} + \gamma \frac{\partial \mathcal{L}}{\partial \mathbf{X}_L^k} \quad (13)$$

$$\mathbf{G}_l^k = (1 - \alpha) \tilde{\mathbf{A}} \mathbf{G}_{l+1} + \alpha \frac{\partial \mathcal{L}}{\partial \mathbf{X}_L^k} \quad \forall l = L - 1, \dots, 0 \quad (14)$$

where  $\frac{\partial \mathcal{L}}{\partial \mathbf{X}_*^k}$  is approximated by  $\frac{\partial \mathcal{L}}{\partial \mathbf{X}_L^k}$ , and  $\mathbf{G}_0$  provides an approximation for gradient  $\frac{\partial \mathcal{L}}{\partial \mathbf{X}_{\text{in}}^k}$  so that the gradient of model parameters, i.e.,  $\frac{\partial \mathcal{L}}{\partial \Theta^k}$ , can be further computed by chain rules. Similar to the forward computation, the momentum correction in Eq. (13) compensates the gradient changes over training iterations by mixing the current gradient  $\frac{\partial \mathcal{L}}{\partial \mathbf{X}_L^k}$  and history gradient  $\frac{\partial \mathcal{L}}{\partial \mathbf{X}_{\text{in}}^{k-1}}$  with hyperparameter  $\gamma$ .

**Remark 3 (Computation and memory efficiency).** *Since LazyGNN uses very few aggregation layers, it significantly reduces the computation and memory cost. Moreover, LazyGNN does not need to store intermediate hidden activation values in the aggregation layers because the computation of implicit gradient is agnostic to the forward propagation trajectory, which further reduces memory cost.*

**Remark 4 (Communication efficiency).** *LazyGNN provides a fundamental algorithmic improvement to significantly reduce the communication cost in cross-device feature aggregations for distributed GNN training because it uses fewer propagations and communication rounds. This is orthogonal and complementary to many existing system-level strategies that mitigate the significant feature communication overhead in distributed GNN training (Tripathy et al., 2020; Wan et al., 2022b;a; Zhang et al., 2023).*

### 3.2. Mini-batch LazyGNN with Subgraph Sampling

As illustrated in Section 3.1, LazyGNN solves the inherent neighborhood explosion problem so that when handling large-scale graphs, the mini-batch sampling only needs to sample neighboring nodes within a few hop distances. To further demonstrate the superior scalability, we introduce a mini-batch LazyGNN that enhances the computation and memory efficiency with efficient data sampling as shown in Figure 4. In each training iteration  $k$ , we sample a target node set  $V_1^k$  and their  $L$ -hop neighbor set  $V_2^k$ , and we denote the union of these two nodes as  $V^k = V_1^k \cup V_2^k$ . An induced subgraph  $\tilde{\mathbf{A}}_{V^k}$  is constructed based on the node set  $V^k$ . Note that LazyGNN works well with small  $L \in \{1, 2\}$  so that the data sampling is extremely efficient.

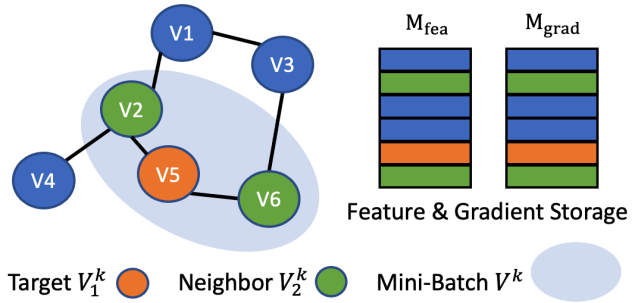


Figure 4. Mini-batch LazyGNN with Feature & Gradient Storage.

**Forward Computation.** The forward computation of mini-batch LazyGNN works as follows:

$$(\mathbf{X}_{\text{in}}^k)_{V_k} = f((\mathbf{X}_{\text{fea}})_{V^k}, \Theta^k), \quad (15)$$

$$(\mathbf{X}_0)_{V_k} = (1 - \beta)(\mathbf{M}_{\text{fea}})_{V^k} + \beta(\mathbf{X}_{\text{in}}^k)_{V_k} \quad (16)$$

$$(\mathbf{X}_{l+1}^k)_{V^k} = (1 - \alpha) \tilde{\mathbf{A}}_{V^k} (\mathbf{X}_l^k)_{V^k} + \alpha (\mathbf{X}_{\text{in}}^k)_{V^k}, \quad \forall l \quad (17)$$

$$(\mathbf{M}_{\text{fea}})_{V_1^k} = (\mathbf{X}_L^k)_{V_1^k} \quad (18)$$

The node features  $(\mathbf{X}_{\text{fea}})_{V^k}$  for the node set  $V^k$  are sampled to form a mini-batch. The corresponding diffused node features  $(\mathbf{M}_{\text{fea}})_{V^k}$  of the same node set are retrieved from the feature storage  $\mathbf{M}_{\text{fea}}$  and then combined with current node features  $(\mathbf{X}_{\text{in}}^k)_{V_k}$  in Eq. (16). The lazy forward propagation runs on the subgraph  $\tilde{\mathbf{A}}_{V^k}$  in Eq. (17). Finally,  $(\mathbf{X}_L^k)_{V_1^k}$  provides the prediction for target nodes  $V_1^k$ , which is further maintained in the features storage  $\mathbf{M}_{\text{fea}}$ . The embeddings of neighbor nodes  $(\mathbf{X}_L^k)_{V_2^k}$  are not accurate and not stored.

**Backward Computation.** The backward computation of mini-batch LazyGNN works as follows:

$$(\mathbf{G}_L^k)_{V^k} = (1 - \gamma)(\mathbf{M}_{\text{grad}})_{V^k} + \gamma \frac{\partial \mathcal{L}}{\partial (\mathbf{X}_L^k)_{V^k}} \quad (19)$$

$$(\mathbf{G}_l^k)_{V^k} = (1 - \alpha) \tilde{\mathbf{A}}_{V^k} (\mathbf{G}_{l+1}^k)_{V^k} + \frac{\alpha}{\partial (\mathbf{X}_L^k)_{V^k}} \frac{\partial \mathcal{L}}{\partial (\mathbf{X}_L^k)_{V^k}} \quad \forall l \quad (20)$$

$$(\mathbf{M}_{\text{grad}})_{V_1^k} = (\mathbf{G}_0^k)_{V_1^k} \quad (21)$$

In Eq. (19),  $(\mathbf{M}_{\text{grad}})_{V^k}$ , the history gradient with respect to node features  $(\mathbf{X}_{\text{in}}^{k-1})_{V^k}$ , is retrieved from the gradient storage  $\mathbf{M}_{\text{grad}}$  and then combined with  $\frac{\partial \mathcal{L}}{\partial (\mathbf{X}_L^k)_{V^k}}$ , the gradient with respect to current node features  $(\mathbf{X}_L^k)_{V^k}$ . The lazy backward propagation runs on the subgraph  $\mathbf{A}_{V^k}$  in Eq. (20). Finally,  $(\mathbf{G}_0^k)_{V_1^k}$  provides an estimation for  $\frac{\partial \mathcal{L}}{\partial (\mathbf{X}_{\text{in}}^k)_{V_1^k}}$ , which is used to compute gradient of parameters  $\Theta^k$  by chain rules and then maintained in the gradient storage  $\mathbf{M}_{\text{grad}}$ . Likewise, the gradients of neighbor nodes  $\frac{\partial \mathcal{L}}{\partial (\mathbf{X}_{\text{in}}^k)_{V_2^k}}$  are not accurate and will not be stored in the gradient memory.

### 3.3. Complexity Analysis

As discussed in Section 3.1 and Section 3.2, LazyGNN is scalable due to its efficiency in computation, memory, data sampling, and communication. In this section, we provide complexity analyses to further demonstrate its superior scalability. Since the major efficiency difference lies in feature aggregation layers, we focus on the complexity of feature aggregations. Suppose  $L$  is the number of propagation layers,  $H$  is the size of hidden units,  $N$  is the number of nodes,  $M$  is the number of edges. For simplicity, we assume that  $H$  is the same for all hidden layers.

**Computation complexity.** Performing one feature aggregation requires a sparse-dense matrix multiplication that needs  $\mathcal{O}(MH)$  operations. Therefore, the computation complexity for forward feature aggregations and backward gradient aggregations is  $\mathcal{O}(2LMH)$  per epoch. However, the number of layers  $L$  in LazyGNN is much smaller than existing approaches, so it significantly reduces the computation cost.

**Memory complexity.**  $\mathcal{O}(NH)$  memory space is required for the storage of history feature  $\mathbf{X}_L^k$  (i.e.,  $\mathbf{M}_{\text{fea}}$ ) and history gradient  $\frac{\partial \mathcal{L}}{\partial \mathbf{X}_{\text{in}}^k}$  (i.e.,  $\mathbf{M}_{\text{grad}}$ ) in LazyGNN. LazyGNN does not store the intermediate state at each feature aggregation layer because the backward gradient computation is agnostic to the forward computation trajectory. Therefore, the total memory complexity for LazyGNN is  $\mathcal{O}(NH)$ , which is independent of the number of aggregation layers. In contrast, existing state-of-art history-embedding based models, such as GNNAutoScale (Fey et al., 2021) and GraphFM (Yu et al., 2022), require the storage of feature embeddings in each layer, which leads to the memory complexity  $\mathcal{O}(LNH)$ . Moreover, they also require the storage of all intermediate layers for gradient computation. Their memory cost increases linearly with the number of layers, which is essential in capturing long-distance dependency in those models. In fact, because of the large storage requirements, GAS chooses to save the history embedding in CPU memory, but the data movement between CPU and GPU can be dominating compared with the movement in GPU as verified in Section 4.2.

**Sampling efficiency.** Mini-batch LazyGNN further reduces the computation and memory cost by data sampling, which enjoys the same benefits as existing sampling-based methods. Compared to classic sampling-based models, such as GraphSAGE and GraphSAINT, mini-batch LazyGNN only needs to sample neighbors in one hop or two hops distance so that the data sampling is efficient. This data sampling efficiency might bring significant practical speedup for large-scale problems where the data loading of node features and subgraphs can be dominating (Ma et al., 2022).

**Communication efficiency.** Recent works (Tripathy et al., 2020; Wan et al., 2022b;a; Zhang et al., 2023) clearly demonstrate that the significant feature communication overhead in propagation layers dominates the running time in the parallel training of GNNs. Specifically, the communication time divided by the total training time is over 80% for ogbn-products as shown in the work (Wan et al., 2022b). Therefore, existing works propose various system-level strategies to mitigate this bottleneck. Although this work does not focus on the distributed setting, it is worth emphasizing that LazyGNN is promising for parallelism since it requires fewer communication rounds due to lazy propagation.

Above analyses indicate that LazyGNN is highly efficient in computation, memory, data sampling, and communication. The practical running time and memory cost will be discussed in Section 4.2.

## 4. Experiments

In this section, we provide comprehensive experiments to validate the superior prediction performance and scalability of LazyGNN. Specifically, we try to answer the following questions: (Q1) Can LazyGNN achieve strong prediction accuracy on large-scale graph benchmarks? (Section 4.1) (Q2) Can LazyGNN handle large graphs more efficiently than existing approaches? (Section 4.2) (Q3) What are the impacts of the hyperparameters in LazyGNN? (Section 4.3)

### 4.1. Prediction Performance

**Experimental settings.** We conduct experiments on multiple large-scale graph datasets including REDDIT, YELP, FLICKR, ogbn-arxiv, and ogbn-products (Hu et al., 2020). We evaluate the graph representation learning by node classification accuracy in the semi-supervised setting. We provide a performance comparison with multiple baselines including GCN (Kipf & Welling, 2016), GraphSage (Hamilton et al., 2017), FastGCN (Chen et al., 2018a), LADIES (Zou et al., 2019), VR-GCN (Chen et al., 2017), MVS-GNN (Cong et al., 2020), Cluster-GCN (Chiang et al., 2019), GraphSAINT (Zeng et al., 2020), SGC (Wu et al., 2019), SIGN (Rossi et al., 2020), GAS (Fey et al., 2021) and VQ-GNN (Ding et al., 2021). The hyperparameter tuning of

baselines closely follows the setting in GNNAutoScale (Fey et al., 2021).

For LazyGNN, hyperparameters are tuned from the following search space: (1) learning rate:  $\{0.01, 0.001, 0.0001\}$ ; (2) weight decay:  $\{0, 5e-4, 5e-5\}$ ; (3) dropout:  $\{0.1, 0.3, 0.5, 0.7\}$ ; (4) propagation layers :  $L \in \{1, 2\}$ ; (5) MLP layers:  $\{3, 4\}$ ; (6) MLP hidden units:  $\{256, 512\}$ ; (7)  $\alpha \in \{0.01, 0.1, 0.2, 0.5, 0.8\}$ ; (8)  $\beta$  and  $\gamma$  are simply set as 0.5 in most cases, but a further tuning can improve the performance. While LazyGNN is compatible with any sampling method, we adopt the subgraph sampling strategy as in GNNAutoScale to ensure a fair comparison.

**Performance analysis.** From the accuracy as summarized in Table 1, we can make the following observations:

- LazyGNN achieves state-of-art performance on almost all datasets. In particular, LazyGNN achieves 73.2% and 82.3% accuracy on ogbn-arxiv and ogbn-products. The only exceptions are that GraphSAINT slightly outperforms LazyGNN by 0.2% on REDDIT and that GCNII outperforms LazyGNN by 1.1% on FLICKR. However, LazyGNN is much more efficient than GraphSAINT and GCNII in computation and memory cost. Note that GCNII achieves the reported performance with 10 transformation layers and 8 propagation layers while LazyGNN only uses 4-layer MLP and 2-layer propagation. We believe a stronger performance can be achieved with a more thorough architecture tuning on LazyGNN.
- Importantly, we observe that full-batch LazyGNN closely matches mini-batch LazyGNN, which indicates that the subgraph sampling in LazyGNN can significantly improve the scalability (as will be discussed in Section 4.2) but do not heavily hurt the prediction performance. Sometimes, LazyGNN even improves the generalization performance due to sampling (such as for REDDIT).
- LazyGNN also consistently outperforms APPNP with 10 propagation layers, which indicates the advantage of capturing long-distance dependency in graphs by lazy propagation. The reproduction of GAS+GCN on YELP is much worse than reported in the work (Fey et al., 2021), therefore we omit the result.

#### 4.2. Efficiency Analysis

To verify the scalability of LazyGNN, we provide empirical efficient analysis compared with state-of-art end-to-end trainable and scalable GNNs such as GNNAutoScale (GAS) since it has been shown to be the most efficient algorithm for training large-scale GNNs (Fey et al., 2021). Specifically, we measure the memory usage and running time for the training on ogbn-products dataset using the same batch size. We run GAS using the authors’ code which stores the

Table 1. Prediction accuracy (%) on large-scale graph datasets. “full” and “mini” stand for full-batch and mini-batch respectively. OOM stands for out-of-memory.

# nodes # edges Method	230K 11.6M REDDIT	89K 450K FLICKR	717K 7.9M YELP	169K 1.2M ogbn- arxiv	2.4M 61.9M ogbn- products
GraphSAGE	95.4	50.1	63.4	71.5	78.7
FastGCN	93.7	50.4	—	—	—
LADIES	92.8	—	—	—	—
VR-GCN	94.5	—	61.5	—	—
MVS-GNN	94.9	—	62.0	—	—
Cluster-GCN	96.6	48.1	60.9	—	79.0
GraphSAINT	<b>97.0</b>	51.1	65.3	—	79.1
SGC	96.4	48.2	64.0	—	—
SIGN	96.8	51.4	63.1	—	77.6
VQ-GNN	94.5	—	—	70.6	—
GCN (full)	95.4	53.7	OOM	71.6	OOM
APPNP (full)	96.1	53.4	OOM	71.8	OOM
GCNII (full)	96.1	55.3	OOM	72.8	OOM
GCN (GAS)	95.4	54.0	—	71.7	76.7
APPNP (GAS)	96.0	52.4	63.8	71.9	76.2
GCNII (GAS)	96.7	<b>55.3</b>	65.1	72.5	77.2
LazyGNN (full)	96.2	54.2	OOM	<b>73.2</b>	OOM
LazyGNN (mini)	96.8	54.0	<b>65.4</b>	73.0	<b>82.3</b>

feature embedding for each layer in CPU memory by default since it requires large memory for embedding storage. We also re-implement GAS by storing all memory on the GPU for a better comparison. For LazyGNN, the memory for storing two small tensors (i.e., feature and gradient) is required. Furthermore, we measure the running time of all models for two cases depending on whether the storage is on CPU or GPU memory. Note that all methods use similar number of training epochs as will be verified in Section 4.3, so we compare them using the running time per epoch. For GAS baselines, we use the architecture that can reach the best performance. From the measurements summarized in Table 2, we can make the following observations:

- LazyGNN (GPU) has a much shorter running time (7.7s) per epoch compared with all baselines. All methods become much slower if the memory needs to be moved between GPU and CPU.
- LazyGNN only requires small memory space and only 878 MB of additional storage is required to store two small tensors for feature and gradient regardless of the number of layers. In contrast, GAS requires large memory storage due to the feature storage for each layer. This verifies that LazyGNN is memory efficient.
- LazyGNN (GPU+CPU) is faster than APPNP (GAS) because APPNP requires more data movement for feature propagation. This further verifies the computation and memory efficiency of LazyGNN. An even better running time of LazyGNN is expected if the data movement is

optimized as the implementation in GAS using asynchronous calls, but we leave it as a future work.

Table 2. Memory usage (MB) and running time (seconds per epoch) on ogbn-products.

Method	Setting	GPU MEMORY	CPU MEMORY	RUNNING TIME
GCN (GAS)		2512	4783	28.8
GCNII (GAS)	GPU+CPU	3035	4783	44.3
APPNP (GAS)		2142	3952	84.5
LazyGNN (mini)		2101	878	44.5
GraphSage		4781	—	33.7
GCN (GAS)		7296	—	11.7
GCNII (GAS)	GPU Only	7820	—	18.2
APPNP (GAS)		6097	—	10.5
LazyGNN (mini)		<b>2981</b>	—	<b>7.7</b>

### 4.3. Ablation Study

We provide detailed ablation studies on the impacts of hyperparameter settings in LazyGNN.

**Oversmoothing.** We measure the L2 distance between connected nodes to show the smoothness of GCN, APPNP, and LazyGNN on Cora. Table 3 shows that the smoothness value of GCN becomes smaller and smaller as the number of layers increases, which means the node features become over-smoothed and indistinguishable. But for LazyGNN and APPNP, when more layers are being stacked, the smoothness values become stable so that the node features can still be distinguished without suffering from over-smoothing. This provides direct evidence that LazyGNN does not suffer from the over-smoothing problem as discussed in Remark 2.

Table 3. L<sub>2</sub> feature distance between connected nodes (smoothness) with a varying number of propagation layers on Cora.

Layers	GCN	APPNP	LAZYGNN
L = 1	336	353	302
L = 5	315	305	285
L = 10	297	296	290
L = 15	263	304	291
L = 20	225	302	292

**Lazy Propagation.** We conduct the ablation study to analyze the impact of feature propagation layers and long-distance dependency. Specifically, we use the same MLP for LazyGNN and APPNP and change the number of propagation layers  $L$ . The comparison in Table 4 shows that:

- The performance of APPNP (GAS) improves with the use of more propagation layers (larger  $L$ ). This trend is pretty clear on ogbn-arxiv, which verifies the benefits of capturing long-distance dependency on graphs. The improvement on ogbn-products is relatively minor due

to the large staleness of feature memory caused by GAS when a smaller batch size has to be used.

- LazyGNN archives great performance with even 1 propagation layer (72.5% for ogbn-arxiv and 79.8% for ogbn-products). It has already achieved state-of-the-art performance (73.0% for ogbn-arxiv and 82.3% for ogbn-products) with 2 propagation layers. These results confirm LazyGNN’s advantages and capability in capturing long-distance dependency in graphs with very few feature propagation layers.

Table 4. Prediction accuracy with a varying number of propagation layers on ogbn-arxiv and ogbn-products.

Method	OGBN-ARXIV (%)	OGBN-PRODUCTS (%)
APPNP (GAS) (L=1)	70.6	75.2
APPNP (GAS) (L=2)	71.4	75.5
APPNP (GAS) (L=3)	71.5	75.7
APPNP (GAS) (L=4)	71.8	75.8
APPNP (GAS) (L=5)	72.2	76.0
APPNP (GAS) (L=8)	72.5	76.1
APPNP (GAS) (L=10)	72.7	76.2
LazyGNN (L=1)	72.5	79.8
LazyGNN (L=2)	<b>73.0</b>	<b>82.3</b>
LazyGNN (L=3)	<b>73.0</b>	OOM

**Batch Size.** We investigate the impact of the batch size on data sampling. Note that we adopt a cluster-based subgraph sampling so that the batch size is determined by the number of clusters since one cluster is sampled in each batch. In other words, more clusters result in a smaller batch size. The comparison in Table 5 demonstrates that a larger batch size (less clusters) brings better performance because the sampling variance is reduced. Overall the performance is quite stable using varying batch sizes (clusters).

Table 5. Prediction accuracy of mini-batch LazyGNN with different batch sizes on ogbn-products.

Clusters	50	100	150	200
Accuracy(%)	82.3	81.8	81.5	80.6

**Sensitivity Analysis.** We provide a detailed sensitivity analysis for the two additional hyperparameters  $\beta$  and  $\gamma$  in LazyGNN. The results in Table 6 show that: (1)  $\{\beta = 0, \gamma = 0\}$  produces the worst performance (71.0%) because it fully trusts the history embedding that might be outdated due to the additional variations caused by dropout as demonstrated in Section 2.2; (2)  $\{\beta = 1, \gamma = 1\}$  performs slightly better (71.7%) even without lazy propagation because 2 propagation layers can provide a good approximate solution; (3)  $\{\beta = 0.5, \gamma = 0.5\}$  performs the best (73.0%) because it achieves a good trade-off between historical information and current iteration. It compensates for the staleness in the history storage as designed; (4) The performance of LazyGNN is quite stable for a large range setting



of  $\beta$  and  $\gamma$ . In fact, we fix  $\beta = \gamma = 0.5$  for LazyGNN in most experiments for simplicity. Therefore, it requires almost negligible effort for additional hyperparameter tuning.

Table 6. Prediction accuracy of mini-batch LazyGNN for different hyperparameter settings on ogbn-arxiv.

Hyperparameter settings	ACCURACY (%)
$\beta = 0.0, \gamma = 0.0$	71.0
$\beta = 1.0, \gamma = 1.0$	71.7
$\beta = 0.5, \gamma = 0.1$	72.5
$\beta = 0.5, \gamma = 0.5$	73.0
$\beta = 0.5, \gamma = 0.9$	72.8
$\beta = 0.1, \gamma = 0.5$	72.4
$\beta = 0.5, \gamma = 0.5$	73.0
$\beta = 0.9, \gamma = 0.5$	72.7

**Convergence.** We show the convergence of LazyGNN during the training process using ogbn-arxiv dataset. The convergence of validation accuracy in Figure 5 demonstrates that LazyGNN has a comparable convergence speed with GCN (GAS) and GCNII (GAS), and is slightly faster than APPNP (GAS) in terms of the number of training epochs.




Figure 5. Validation performance versus training epochs.

## 5. Related Work

It has been generally demonstrated that it is beneficial to capture long-distance relations in graphs by stacking more feature aggregation layers or unrolling various fixed point iterations in GNNs (Gasteiger et al., 2018; Gu et al., 2020; Liu et al., 2020; Chen et al., 2020a; Li et al., 2021; Ma et al., 2020; Pan et al., 2020; Zhu et al., 2021; Chen et al., 2020b). But these works suffer from scalability concerns due to the neighborhood explosion problem (Hamilton et al., 2017).

A large body of existing research focuses on improving the efficiency and scalability of large-scale GNNs using various novel designs, such as sampling methods, pre-computing or post-computing methods, and distributed methods. Sampling methods adopt mini-batch training strategies to reduce

computation and memory requirements by sampling nodes and edges. They mitigate the neighbor explosion issue by either removing neighbors (Hamilton et al., 2017; Chen et al., 2018a; Zeng et al., 2020; Zou et al., 2019) or updating with feature memory (Fey et al., 2021; Yu et al., 2022). Pre-computing or post-computing methods separate the feature aggregation and prediction models into two stages, such as pre-computing the feature aggregation before training (Wu et al., 2019; Rossi et al., 2020; Sun et al., 2021; Zhang et al., 2022; Bojchevski et al., 2020) or post-processing with label propagation after training (Zhu, 2005; Huang et al., 2020). Distributed methods distribute large graphs to multiple servers and parallelize GNNs training (Chiang et al., 2019; Chai et al., 2022; Shao et al., 2022). In retrospect, these existing approaches still suffer from various limitations, such as high costs in computation, memory, and communication as well as performance degradation due to large approximation errors or multi-stage training.

Different from existing works, in this work, we propose LazyGNN from a substantially different and novel perspective and propose to capture the long-distance dependency in graphs by shallower models instead of deeper models. This leads to a much more efficient LazyGNN for graph representation learning. Moreover, existing approaches for scalable GNNs can be used to further accelerate LazyGNN. The proposed mini-batch LazyGNN is a promising example.

LazyGNN also draws insight from existing research on graph signal processing and implicit modeling. The forward construction of LazyGNN follows the optimization perspective of GNNs (Ma et al., 2020), and it can be easily generalized to other fixed point iterations with various diffusion properties. LazyGNN is related to recent works in implicit modeling, such as Neural Ordinary Differential Equations (Chen et al., 2018b), Implicit Deep Learning (El Ghaoui et al., 2021), Deep Equilibrium Models (Bai et al., 2019), and Implicit GNNs (Gu et al., 2020). LazyGNN focuses on developing shallow models with highly scalable and efficient computation while these implicit models demand significantly more computation resources.

## 6. Conclusions

We propose LazyGNN, a novel shallow model to solve the neighborhood explosion problem in large-scale GNNs while capturing long-distance dependency through lazy propagation. We also develop a highly scalable and efficient variant, mini-batch LazyGNN, to handle large graphs. Comprehensive experiments demonstrate its superior prediction performance and efficiency on large-scale problems. The proposed lazy propagation provides a promising algorithmic strategy complementary to existing efforts. We plan to explore its application on other types of GNN models and its further acceleration using other scalable techniques in the future.

## References

- Bai, S., Kolter, J. Z., and Koltun, V. Deep equilibrium models. *Advances in Neural Information Processing Systems*, 32, 2019.
- Bojchevski, A., Klicpera, J., Perozzi, B., Kapoor, A., Blais, M., Róžemberczki, B., Lukasiak, M., and Günnemann, S. Scaling graph neural networks with approximate pagerank. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 2464–2473, 2020.
- Chai, Z., Bai, G., Zhao, L., and Cheng, Y. Distributed graph neural network training with periodic historical embedding synchronization. *arXiv preprint arXiv:2206.00057*, 2022.
- Chen, J., Zhu, J., and Song, L. Stochastic training of graph convolutional networks with variance reduction. *arXiv preprint arXiv:1710.10568*, 2017.
- Chen, J., Ma, T., and Xiao, C. Fastgcn: fast learning with graph convolutional networks via importance sampling. *arXiv preprint arXiv:1801.10247*, 2018a.
- Chen, M., Wei, Z., Huang, Z., Ding, B., and Li, Y. Simple and deep graph convolutional networks. In *International Conference on Machine Learning*, pp. 1725–1735. PMLR, 2020a.
- Chen, Q., Wang, Y., Wang, Y., Yang, J., and Lin, Z. Optimization-induced graph implicit nonlinear diffusion. In *International Conference on Machine Learning*, pp. 3648–3661. PMLR, 2022.
- Chen, R. T., Rubanova, Y., Bettencourt, J., and Duvenaud, D. K. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018b.
- Chen, S., Eldar, Y. C., and Zhao, L. Graph unrolling networks: Interpretable neural networks for graph signal denoising, 2020b.
- Chiang, W.-L., Liu, X., Si, S., Li, Y., Bengio, S., and Hsieh, C.-J. 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*, pp. 257–266, 2019.
- Cong, W., Forsati, R., Kandemir, M., and Mahdavi, M. Minimal variance sampling with provable guarantees for fast training of graph neural networks. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 1393–1403, 2020.
- Ding, M., Kong, K., Li, J., Zhu, C., Dickerson, J., Huang, F., and Goldstein, T. Vq-gnn: A universal framework to scale up graph neural networks using vector quantization. *Advances in Neural Information Processing Systems*, 34: 6733–6746, 2021.
- El Ghaoui, L., Gu, F., Travacca, B., Askari, A., and Tsai, A. Implicit deep learning. *SIAM Journal on Mathematics of Data Science*, 3(3):930–958, 2021.
- Fey, M., Lenssen, J. E., Weichert, F., and Leskovec, J. Gnnautoscale: Scalable and expressive graph neural networks via historical embeddings. In *International Conference on Machine Learning*, pp. 3294–3304. PMLR, 2021.
- Gasteiger, J., Bojchevski, A., and Günnemann, S. Predict then propagate: Graph neural networks meet personalized pagerank. *arXiv preprint arXiv:1810.05997*, 2018.
- Gasteiger, J., Bojchevski, A., and Günnemann, S. Combining neural networks with personalized pagerank for classification on graphs. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=H1gL-2A9Ym>.
- Gu, F., Chang, H., Zhu, W., Sojoudi, S., and El Ghaoui, L. Implicit graph neural networks. *Advances in Neural Information Processing Systems*, 33:11984–11995, 2020.
- Hamilton, W., Ying, Z., and Leskovec, J. Inductive representation learning on large graphs. *Advances in neural information processing systems*, 30, 2017.
- Hamilton, W. L. Graph representation learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 14(3):1–159, 2020.
- Hochreiter, S. and Schmidhuber, J. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- Hu, W., Fey, M., Zitnik, M., Dong, Y., Ren, H., Liu, B., Catasta, M., and Leskovec, J. Open graph benchmark: Datasets for machine learning on graphs. *Advances in neural information processing systems*, 33:22118–22133, 2020.
- Huang, Q., He, H., Singh, A., Lim, S.-N., and Benson, A. R. Combining label propagation and simple models out-performs graph neural networks. *arXiv preprint arXiv:2010.13993*, 2020.
- Jia, J. and Benson, A. R. A unifying generative model for graph learning algorithms: Label propagation, graph convolutions, and combinations. *SIAM Journal on Mathematics of Data Science*, 4(1):100–125, 2022.

- Jiang, Z., Han, X., Fan, C., Liu, Z., Zou, N., Mostafavi, A., and Hu, X. Fmp: Toward fair graph message passing against topology bias. *arXiv preprint arXiv:2202.04187*, 2022.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Kipf, T. N. and Welling, M. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.
- Klicpera, J., Bojchevski, A., and Günnemann, S. Predict then propagate: Graph neural networks meet personalized pagerank. *arXiv preprint arXiv:1810.05997*, 2018.
- Li, G., Müller, M., Ghanem, B., and Koltun, V. Training graph neural networks with 1000 layers. In *International conference on machine learning*, pp. 6437–6449. PMLR, 2021.
- Liu, M., Gao, H., and Ji, S. Towards deeper graph neural networks. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*, pp. 338–348, 2020.
- Liu, X., Ding, J., Jin, W., Xu, H., Ma, Y., Liu, Z., and Tang, J. Graph neural networks with adaptive residual. *Advances in Neural Information Processing Systems*, 34: 9720–9733, 2021a.
- Liu, X., Jin, W., Ma, Y., Li, Y., Liu, H., Wang, Y., Yan, M., and Tang, J. Elastic graph neural networks. In *International Conference on Machine Learning*, pp. 6837–6849. PMLR, 2021b.
- Ma, Y. and Tang, J. *Deep learning on graphs*. Cambridge University Press, 2021.
- Ma, Y., Liu, X., Zhao, T., Liu, Y., Tang, J., and Shah, N. A unified view on graph neural networks as graph signal denoising. *arXiv preprint arXiv:2010.01777*, 2020.
- Ma, Y., Gong, P., Yi, J., Yao, Z., Wang, M., Li, C., He, Y., and Yan, F. Bifeat: Supercharge gnn training via graph feature quantization. *arXiv preprint arXiv:2207.14696*, 2022.
- Pan, X., Shiji, S., and Gao, H. A unified framework for convolution-based graph neural networks. <https://openreview.net/forum?id=zUMD-Fb9Bt>, 2020.
- Rossi, E., Frasca, F., Chamberlain, B., Eynard, D., Bronstein, M., and Monti, F. Sign: Scalable inception graph neural networks. *arXiv preprint arXiv:2004.11198*, 7:15, 2020.
- Shao, Y., Li, H., Gu, X., Yin, H., Li, Y., Miao, X., Zhang, W., Cui, B., and Chen, L. Distributed graph neural network training: A survey. *arXiv preprint arXiv:2211.00216*, 2022.
- Sun, C., Gu, H., and Hu, J. Scalable and adaptive graph neural networks with self-label-enhanced training. *arXiv preprint arXiv:2104.09376*, 2021.
- Tripathy, A., Yelick, K., and Buluç, A. Reducing communication in graph neural network training. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, SC '20*. IEEE Press, 2020. ISBN 9781728199986.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- Veličković, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., and Bengio, Y. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017.
- Wan, C., Li, Y., Li, A., Kim, N. S., and Lin, Y. Bns-gcn: Efficient full-graph training of graph convolutional networks with partition-parallelism and random boundary node sampling. *Proceedings of Machine Learning and Systems*, 4:673–693, 2022a.
- Wan, C., Li, Y., Wolfe, C. R., Kyrillidis, A., Kim, N. S., and Lin, Y. PipeGCN: Efficient full-graph training of graph convolutional networks with pipelined feature communication. In *International Conference on Learning Representations*, 2022b. URL <https://openreview.net/forum?id=kSwqMH0zn1F>.
- Wu, F., Souza, A., Zhang, T., Fifty, C., Yu, T., and Weinberger, K. Simplifying graph convolutional networks. In *International conference on machine learning*, pp. 6861–6871. PMLR, 2019.
- Yang, Y., Liu, T., Wang, Y., Zhou, J., Gan, Q., Wei, Z., Zhang, Z., Huang, Z., and Wipf, D. Graph neural networks inspired by classical iterative algorithms. In *International Conference on Machine Learning*, pp. 11773–11783. PMLR, 2021a.
- Yang, Y., Wang, Y., Huang, Z., and Wipf, D. Implicit vs unfolded graph neural networks. *arXiv preprint arXiv:2111.06592*, 2021b.
- Ying, R., He, R., Chen, K., Eksombatchai, P., Hamilton, W. L., and Leskovec, J. Graph convolutional neural networks for web-scale recommender systems. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, pp. 974–983, 2018.

- Yu, H., Wang, L., Wang, B., Liu, M., Yang, T., and Ji, S. Graphfm: Improving large-scale gnn training via feature momentum. In *International Conference on Machine Learning*, pp. 25684–25701. PMLR, 2022.
- Zeng, H., Zhou, H., Srivastava, A., Kannan, R., and Prasanna, V. Graphsaint: Graph sampling based inductive learning method. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=BJe8pkHFwS>.
- Zhang, M., Hu, Q., Sun, P., Wen, Y., and Zhang, T. Boosting distributed full-graph gnn training with asynchronous one-bit communication. *arXiv preprint arXiv:2303.01277*, 2023.
- Zhang, W., Yin, Z., Sheng, Z., Li, Y., Ouyang, W., Li, X., Tao, Y., Yang, Z., and Cui, B. Graph attention multi-layer perceptron. *arXiv preprint arXiv:2206.04355*, 2022.
- Zhou, D., Bousquet, O., Lal, T., Weston, J., and Schölkopf, B. Learning with local and global consistency. *Advances in neural information processing systems*, 16, 2003.
- Zhu, M., Wang, X., Shi, C., Ji, H., and Cui, P. Interpreting and unifying graph neural networks with an optimization framework, 2021.
- Zhu, X. *Semi-supervised learning with graphs*. Carnegie Mellon University, 2005.
- Zou, D., Hu, Z., Wang, Y., Jiang, S., Sun, Y., and Gu, Q. Layer-dependent importance sampling for training deep and large graph convolutional networks. *Advances in neural information processing systems*, 32, 2019.