
Stochastic Training of Graph Convolutional Networks with Variance Reduction

Jianfei Chen¹ Jun Zhu¹ Le Song^{2,3}

Abstract

Graph convolutional networks (GCNs) are powerful deep neural networks for graph-structured data. However, GCN computes the representation of a node recursively from its neighbors, making the receptive field size grow exponentially with the number of layers. Previous attempts on reducing the receptive field size by subsampling neighbors do not have convergence guarantee, and their receptive field size per node is still in the order of hundreds. In this paper, we develop control variate based algorithms with new theoretical guarantee to converge to a local optimum of GCN regardless of the neighbor sampling size. Empirical results show that our algorithms enjoy similar convergence rate and model quality with the exact algorithm using only two neighbors per node. The running time of our algorithms on a large Reddit dataset is only one seventh of previous neighbor sampling algorithms.

1. Introduction

Graph convolution networks (GCNs) (Kipf & Welling, 2017) generalize convolutional neural networks (CNNs) (LeCun et al., 1995) to graph structured data. The “graph convolution” operation applies same linear transformation to all the neighbors of a node, followed by mean pooling and nonlinearity. By stacking multiple graph convolution layers, GCNs can learn node representations by utilizing information from distant neighbors. GCNs and their variants (Hamilton et al., 2017a; Veličković et al., 2018) have been applied to semi-supervised node classification (Kipf & Welling, 2017), inductive node embedding (Hamilton et al., 2017a), link prediction (Kipf & Welling, 2016; Berg et al., 2017) and knowledge graphs (Schlichtkrull et al., 2017), outperforming multi-layer perceptron (MLP) models that

do not use the graph structure, and graph embedding approaches (Perozzi et al., 2014; Tang et al., 2015; Grover & Leskovec, 2016) that do not use node features.

However, the graph convolution operation makes GCNs difficult to be trained efficiently. The representation of a node at layer L is computed recursively by the representations of all its neighbors at layer $L - 1$. Therefore, the receptive field of a single node grows exponentially with respect to the number of layers, as illustrated in Fig. 1(a), so exactly computing the stochastic gradient is expensive even for a single node. Due to the large receptive field size, Kipf & Welling (2017) propose to train GCN by a batch algorithm, which computes the representations of all the nodes altogether. However, batch algorithms cannot handle large-scale datasets because of their slow convergence and the requirement to fit the entire dataset in GPU memory.

Hamilton et al. (2017a) make an initial attempt to develop stochastic training algorithms for GCNs via a scheme of neighbor sampling (NS). Instead of considering all the neighbors, they randomly subsample $D^{(l)}$ neighbors at the l -th layer. Therefore, they reduce the receptive field size to $\prod_l D^{(l)}$, as shown in Fig. 1(b). They find that for two-layer GCNs, keeping $D^{(1)} = 10$ and $D^{(2)} = 25$ neighbors can achieve comparable performance with the original model. However, there is no theoretical guarantee on the convergence of the stochastic training algorithm with NS. Moreover, the time complexity of NS is still $D^{(1)}D^{(2)} = 250$ times larger than training an MLP, which is unsatisfactory.

In this paper, we develop novel control variate-based stochastic approximation algorithms for GCN by utilizing the historical activations of nodes as a control variate. Our algorithms have new theoretical results on (1) variance reduction from the magnitude of the activation to the magnitude of the difference between current-and-historical activations; (2) exact (zero-variance) predictions at testing time; (3) convergence to a local optimum of GCN during training *regardless of the neighbor sampling size* $D^{(l)}$, with an asymptotically unbiased stochastic gradient. The theoretical properties allow us to significantly reduce the time complexity of stochastic training by sampling only $D^{(l)} = 2$ neighbors per node, yet still retain the quality of the model.

We empirically test our algorithms on six graph datasets, and the results match with the theory. Comparing with NS, our

¹Dept. of Comp. Sci. & Tech., BNRist Center, State Key Lab for Intell. Tech. & Sys., THBI Lab, Tsinghua University, Beijing, 100084, China ²Georgia Institute of Technology ³Ant Financial. Correspondence to: Jun Zhu <dczsj@mail.tsinghua.edu.cn>.

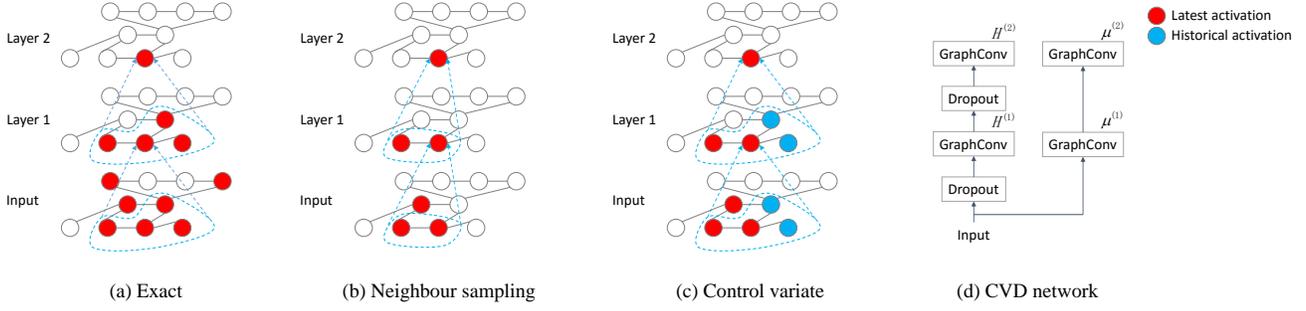


Figure 1. Two-layer graph convolutional networks, and the receptive field of a single vertex.

Dataset	V	E	Degree	Degree 2
Citeseer	3,327	12,431	4	15
Cora	2,708	13,264	5	37
PubMed	19,717	108,365	6	60
NELL	65,755	318,135	5	1,597
PPI	14,755	458,973	31	970
Reddit	232,965	23,446,803	101	10,858

Table 1. Number of vertices, edges, and average number of 1-hop and 2-hop neighbors per node for each dataset. Undirected edges are counted twice and self-loops are counted once.

algorithms significantly reduce the bias and variance of the gradient. Comparing with the exact algorithm which considers all the neighbors, our algorithms with only $D^{(l)} = 2$ neighbors still get the same accuracy at testing time, and achieve similar predictive performance during training *in a comparable number of epochs*, with a much lower time complexity, while these results are not achievable by NS. On the largest Reddit dataset, the training time of our algorithm is 7 times shorter than that of the best-performing competitor among exact, neighbor sampling and importance sampling (Chen et al., 2018) algorithms.

2. Backgrounds

We briefly review graph convolutional networks (GCNs), stochastic training, neighbor sampling, and importance sampling in this section.

2.1. Graph Convolutional Networks

We present our algorithm with a GCN for semi-supervised node classification (Kipf & Welling, 2017). However, the algorithm is neither limited to the task nor the model. Our algorithm is applicable to other models including GraphSAGE-mean (Hamilton et al., 2017a) and graph attention networks (GAT) (Veličković et al., 2018), and other tasks (Kipf & Welling, 2016; Berg et al., 2017; Schlichtkrull et al., 2017; Hamilton et al., 2017b), as long as the model aggregates neighbor activations by averaging.

In the node classification task, we have an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with $V = |\mathcal{V}|$ vertices and $E = |\mathcal{E}|$ edges,

where each vertex v consists of a feature vector x_v and a label y_v . We observe the labels for some vertices $\mathcal{V}_{\mathcal{L}}$. The goal is to predict the labels for the rest vertices $\mathcal{V}_{\mathcal{U}} := \mathcal{V} \setminus \mathcal{V}_{\mathcal{L}}$. The edges are represented as a symmetric $V \times V$ adjacency matrix A , where A_{uv} is the weight of the edge between u and v , and the propagation matrix P is a normalized version of A : $\tilde{A} = A + I$, $\tilde{D}_{uu} = \sum_v \tilde{A}_{uv}$, and $P = \tilde{D}^{-\frac{1}{2}} \tilde{A} \tilde{D}^{-\frac{1}{2}}$. A graph convolution layer is defined as

$$Z^{(l+1)} = PH^{(l)}W^{(l)}, \quad H^{(l+1)} = \sigma(Z^{l+1}), \quad (1)$$

where $H^{(l)}$ is the activation matrix in the l -th layer, whose each row is the activation of a graph node. $H^{(0)} = X$ is the input feature matrix, $W^{(l)}$ is a trainable weight matrix, and $\sigma(\cdot)$ is an activation function. Denote $|\cdot|$ as the cardinality of a set. The training loss is defined as

$$\mathcal{L} = \frac{1}{|\mathcal{V}_{\mathcal{L}}|} \sum_{v \in \mathcal{V}_{\mathcal{L}}} f(y_v, z_v^{(L)}), \quad (2)$$

where $f(\cdot, \cdot)$ is a loss function. A graph convolution layer propagates information to nodes from their neighbors by computing the neighbor averaging $PH^{(l)}$. Let $\mathbf{n}(u)$ be the set of neighbors of node u , and $n(u)$ be its cardinality. The neighbor averaging of node u , $(PH^{(l)})_u = \sum_{v=1}^V P_{uv}h_v^{(l)} = \sum_{v \in \mathbf{n}(u)} P_{uv}h_v^{(l)}$, is a weighted sum of neighbors' activations. Then, a fully-connected layer is applied on all the nodes, with a shared weight matrix $W^{(l)}$ across all the nodes.

We denote the *receptive field* of a node u as all the activations $h_v^{(l)}$ on layer l needed for computing $z_u^{(L)}$. If the layer l is not explicitly mentioned, it is the input layer 0. Intuitively, the receptive field of node u is just all its L -hop neighbors, i.e., nodes that are reachable from u within L hops, as illustrated in Fig. 1(a). When $P = I$, GCN reduces to a multi-layer perceptron (MLP) model which does not use the graph structure. For MLP, the receptive field of a node u is just the node itself.

2.2. Stochastic Training

It is generally expensive to compute the batch gradient $\nabla \mathcal{L} = \frac{1}{|\mathcal{V}_{\mathcal{L}}|} \sum_{v \in \mathcal{V}_{\mathcal{L}}} \nabla f(y_v, z_v^{(L)})$, which involves iterat-

ing over the entire labeled set of nodes. A possible solution is to approximate the batch gradient by a stochastic gradient

$$\frac{1}{|\mathcal{V}_B|} \sum_{v \in \mathcal{V}_B} \nabla f(y_v, z_v^{(L)}), \quad (3)$$

where $\mathcal{V}_B \subset \mathcal{V}_L$ is a minibatch of labeled nodes. However, this gradient is still expensive to compute, due to the large receptive field size. For instance, as shown in Table 1, the number of 2-hop neighbors on the NELL dataset is averagely 1,597, which means in a 2-layer GCN, computing the gradient even for a single node needs $1,597/65,755 \approx 2.4\%$ nodes of the entire graph.

In subsequent sections, two other stochasticity will be introduced besides the random selection of the minibatch: the random sampling of neighbors (Sec. 2.3) and the random dropout of features (Sec. 5).

2.3. Neighbor Sampling

To reduce the receptive field size, Hamilton et al. (2017a) propose a neighbor sampling (NS) algorithm. NS randomly chooses $D^{(l)}$ neighbors for each node at layer l and develops an estimator $\text{NS}_u^{(l)}$ of $(PH^{(l)})_u$ based on Monte-Carlo approximation:

$$(PH^{(l)})_u \approx \text{NS}_u^{(l)} := \frac{n(u)}{D^{(l)}} \sum_{v \in \hat{\mathbf{n}}^{(l)}(u)} P_{uv} h_v^{(l)},$$

where $\hat{\mathbf{n}}^{(l)}(u) \subset \mathbf{n}(u)$ is a subset of $D^{(l)}$ random neighbors. Therefore, NS reduces the receptive field size from all the L -hop neighbors to the number of sampled neighbors, $\prod_{l=1}^L D^{(l)}$. We refer $\text{NS}_u^{(l)}$ as the NS estimator of $(PH^{(l)})_u$, and $(PH^{(l)})_u$ itself as the exact estimator.

Neighbor sampling can also be written in a matrix form as

$$Z^{(l+1)} = \hat{P}^{(l)} H^{(l)} W^{(l)}, \quad H^{(l+1)} = \sigma(Z^{(l+1)}), \quad (4)$$

where the propagation matrix P is replaced by a sparser unbiased estimator $\hat{P}^{(l)}$, i.e., $\mathbb{E} \hat{P}^{(l)} = P$, where $\hat{P}_{uv}^{(l)} = \frac{n(u)}{D^{(l)}} P_{uv}$ if $v \in \hat{\mathbf{n}}^{(l)}(u)$, and $\hat{P}_{uv}^{(l)} = 0$ otherwise. Hamilton et al. (2017a) propose to perform an approximate forward propagation as Eq. (4), and do stochastic gradient descent (SGD) with the auto-differentiation gradient. The approximated gradient has two sources of randomness: the random selection of minibatch $\mathcal{V}_B \subset \mathcal{V}_L$, and the random selection of neighbors.

Though $\hat{P}^{(l)}$ is an unbiased estimator of P , $\sigma(\hat{P}^{(l)} H^{(l)} W^{(l)})$ is *not* an unbiased estimator of $\sigma(PH^{(l)} W^{(l)})$, due to the non-linearity of $\sigma(\cdot)$. In the sequel, both the prediction $Z^{(L)}$ and gradient $\nabla f(y_v, z_v^{(L)})$ obtained by NS are biased, and the convergence of SGD is not guaranteed, unless the sample size $D^{(l)}$ goes to infinity. Because of the biased gradient, the sample size $D^{(l)}$ needs to be large for NS, to keep comparable predictive performance with the exact algorithm. Hamilton

et al. (2017a) choose $D^{(1)} = 10$ and $D^{(2)} = 25$, and the receptive field size $D^{(1)} \times D^{(2)} = 250$ is much larger than one, so the training is still expensive.

2.4. Importance Sampling

FastGCN (Chen et al., 2018) is another sampling-based algorithm similar as NS. Instead of sampling neighbors for each node, FastGCN directly subsample the receptive field for each layer altogether. Formally, it approximates $(PH^{(l)})_u$ with S samples $v_1, \dots, v_S \in \mathcal{V}$ as

$$(PH^{(l)})_u = V \sum_{v=1}^V \frac{1}{V} P_{uv} h_v^{(l)} \approx \frac{V}{S} \sum_{v_s \sim q(v)} P_{uv} h_{v_s}^{(l)} / q(v_s),$$

where they define the importance distribution $q(v) \propto \sum_{u=1}^V P_{uv}^2$. According to the definition of P in Sec. 2.1, we have $q(v) \propto \frac{1}{n(v)} \sum_{(u,v) \in \mathcal{E}} \frac{1}{n(u)}$. We refer to this estimator as importance sampling (IS). Chen et al. (2018) show that IS performs better than using a uniform sample distribution $q(v) \propto 1$. NS can be viewed as an IS estimator with the importance distribution $q(v) \propto \sum_{(u,v) \in \mathcal{E}} \frac{1}{n(u)}$, because each node u has probability $\frac{1}{n(u)}$ to choose the neighbor v . Though IS may have a smaller variance than NS, it still only guarantees the convergence as the sample size S goes to infinity. Empirically, we find IS to work even *worse* than NS because sometimes it can select many neighbors for one node, and no neighbor for another, in which case the activation of the latter node is just meaningless zero.

3. Control Variate Based Algorithm

We present a novel control variate based algorithm that utilizes historical activations to reduce the estimator variance.

3.1. Control Variate Based Estimator

While computing the neighbor average $\sum_{v \in \mathbf{n}(u)} P_{uv} h_v^{(l)}$, we cannot afford to evaluate all the $h_v^{(l)}$ terms because they need to be computed recursively, i.e., we again need the activations $h_w^{(l-1)}$ of all of v 's neighbors w .

Our idea is to maintain the history $\bar{h}_v^{(l)}$ for each $h_v^{(l)}$ as an affordable approximation. Each time when $h_v^{(l)}$ is computed, we update $\bar{h}_v^{(l)}$ with $h_v^{(l)}$. We expect $\bar{h}_v^{(l)}$ and $h_v^{(l)}$ to be similar if the model weights do not change too fast during the training. Formally, let $\Delta h_v^{(l)} = h_v^{(l)} - \bar{h}_v^{(l)}$, we approximate

$$\begin{aligned} (PH^{(l)})_u &= \sum_{v \in \mathbf{n}(u)} P_{uv} \Delta h_v^{(l)} + \sum_{v \in \mathbf{n}(u)} P_{uv} \bar{h}_v^{(l)} \approx \text{CV}_u^{(l)} \\ &:= \frac{n(u)}{D^{(l)}} \sum_{v \in \hat{\mathbf{n}}^{(l)}(u)} P_{uv} \Delta h_v^{(l)} + \sum_{v \in \mathbf{n}(u)} P_{uv} \bar{h}_v^{(l)}, \end{aligned} \quad (5)$$

where we represent $h_v^{(l)}$ as the sum of $\Delta h_v^{(l)}$ and $\bar{h}_v^{(l)}$, and

we only apply Monte-Carlo approximation on the $\Delta h_v^{(l)}$ term. Averaging over all the $\bar{h}_v^{(l)}$ terms is still affordable because they do not need to be computed recursively. Since we expect $h_v^{(l)}$ and $\bar{h}_v^{(l)}$ to be close, Δh_v will be small and $\text{CV}_u^{(l)}$ should have a smaller variance than $\text{NS}_u^{(l)}$. Particularly, if the model weight is kept fixed, $\bar{h}_v^{(l)}$ should eventually equal with $h_v^{(l)}$, so that $\text{CV}_u^{(l)} = 0 + \sum_{v \in \mathbf{n}(u)} P_{uv} \bar{h}_v^{(l)} = \sum_{v \in \mathbf{n}(u)} P_{uv} h_v^{(l)} = (PH^{(l)})_u$, i.e., the estimator has zero variance. This estimator is referred as CV. We will compare the variance of NS and CV estimators in Sec. 3.2 and show that the variance of CV will be eventually zero during the training in Sec. 4. The term $\text{CV}_u^{(l)} - \text{NS}_u^{(l)} = \sum_{v \in \mathbf{n}(u)} P_{uv} \bar{h}_v^{(l)} - \frac{n(u)}{D^{(l)}} \sum_{v \in \hat{\mathbf{n}}^{(l)}(u)} P_{uv} \bar{h}_v^{(l)}$ is a *control variate* (Ripley, 2009, Chapter 5) added to the neighbor sampling estimator $\text{NS}_u^{(l)}$, to reduce its variance.

In matrix form, let $\bar{H}^{(l)}$ be the matrix formed by stacking $\bar{h}_v^{(l)}$, then CV can be written as

$$Z^{(l+1)} = \left(\hat{P}^{(l)}(H^{(l)} - \bar{H}^{(l)}) + P\bar{H}^{(l)} \right) W^{(l)}. \quad (6)$$

3.2. Variance Analysis

We analyze the variance of the estimators assuming all the features are 1-dimensional. The analysis can be extended to multiple dimensions by treating each dimension separately. We further assume that $\hat{\mathbf{n}}^{(l)}(u)$ is created by sampling $D^{(l)}$ neighbors without replacement from $\mathbf{n}(u)$. The following proposition is proven in Appendix A:

Proposition 1. *If $\hat{\mathbf{n}}^{(l)}(u)$ contains $D^{(l)}$ samples from $\mathbf{n}(u)$ without replacement, then $\text{Var}_{\hat{\mathbf{n}}^{(l)}(u)} \left[\frac{n(u)}{D^{(l)}} \sum_{v \in \hat{\mathbf{n}}^{(l)}(u)} x_v \right] = \frac{C_u^{(l)}}{2D^{(l)}} \sum_{v_1 \in \mathbf{n}(u)} \sum_{v_2 \in \mathbf{n}(u)} (x_{v_1} - x_{v_2})^2$, where $C_u^{(l)} = 1 - (D^{(l)} - 1)/(n(u) - 1)$.*

By Proposition 1, we have $\text{Var}_{\hat{\mathbf{n}}^{(l)}(u)} \left[\text{NS}_u^{(l)} \right] = \frac{C_u^{(l)}}{2D^{(l)}} \sum_{v_1 \in \mathbf{n}(u)} \sum_{v_2 \in \mathbf{n}(u)} (P_{uv_1} h_{v_1}^{(l)} - P_{uv_2} h_{v_2}^{(l)})^2$, in contrast, the variance of the CV estimator is $\text{Var}_{\hat{\mathbf{n}}^{(l)}(u)} \left[\text{CV}_u^{(l)} \right] = \frac{C_u^{(l)}}{2D^{(l)}} \sum_{v_1 \in \mathbf{n}(u)} \sum_{v_2 \in \mathbf{n}(u)} (P_{uv_1} \Delta h_{v_1}^{(l)} - P_{uv_2} \Delta h_{v_2}^{(l)})^2$, which replaces $h_v^{(l)}$ by $\Delta h_v^{(l)}$. Since $\Delta h_v^{(l)}$ is usually much smaller than $h_v^{(l)}$, the CV estimator enjoys much smaller variance than the NS estimator. Furthermore, as we will show in Sec. 4.2, $\Delta h_v^{(l)}$ converges to zero during training, so we achieve not only *variance reduction* but *variance elimination*, as the variance vanishes eventually.

3.3. Implementation Details

Training with the CV estimator is similar as with the NS estimator (Hamilton et al., 2017a). Particularly, each iteration of the algorithm involves the following steps:

Stochastic GCN with Variance Reduction

1. Randomly select a minibatch $\mathcal{V}_B \subset \mathcal{V}_L$ of nodes;
2. Build a computation graph that only contains the activations $h_v^{(l)}$ and $\bar{h}_v^{(l)}$ needed for the current minibatch;
3. Get the predictions by forward propagation as Eq. (6);
4. Get the gradients by backward propagation, and update the parameters by SGD;
5. Update the historical activations.

Step 3 and 4 are handled automatically by frameworks such as TensorFlow (Abadi et al., 2016). The computational graph at Step 2 is defined by the receptive field $\mathbf{r}^{(l)}$ and the propagation matrices $\hat{P}^{(l)}$ at each layer. The receptive field $\mathbf{r}^{(l)}$ specifies the activations $h_v^{(l)}$ of which nodes should be computed for the current minibatch, according to Eq. (6). We can construct $\mathbf{r}^{(l)}$ and $\hat{P}^{(l)}$ from top to bottom, by randomly adding $D^{(l)}$ neighbors for each node in $\mathbf{r}^{(l+1)}$, starting with $\mathbf{r}^{(L)} = \mathcal{V}_B$. We assume $h_v^{(l)}$ is always needed to compute $h_v^{(l+1)}$, i.e., v is always selected as a neighbor of itself. The receptive fields are illustrated in Fig. 1(c), where red nodes are in receptive fields, whose activations $h_v^{(l)}$ are needed, and the histories $\bar{h}_v^{(l)}$ of blue nodes are also needed. Finally, in Step 5, we update $\bar{h}_v^{(l)}$ with $h_v^{(l)}$ for each $v \in \mathbf{r}^{(l)}$. We have the pseudocode for the training in Appendix D.

3.4. Time and Space Complexity

GCN has two main types of computation, namely, the sparse-dense matrix multiplication (SPMM) such as $PH^{(l)}$, and the dense-dense matrix multiplication (GEMM) such as $UW^{(l)}$. We assume that the input node feature is K -dimensional and the first hidden layer is A -dimensional.

For batch GCN, the time complexity is $O(EK)$ for SPMM and $O(VKA)$ for GEMM. For our stochastic training algorithm with control variates, the dominant SPMM computation is the average of neighbor history $P\bar{H}^{(l)}$ for the nodes in $\mathbf{r}^{(l)}$, whose size is $O(|\mathcal{V}_B| \prod_{l=2}^L D^{(l)})$, and each node costs $O(DK)$, where D is the average node degree. Therefore, the time complexity of SPMM is approximately $O(EK \prod_{l=2}^L D^{(l)})$ per epoch. The dominant GEMM computation is the first fully-connected layer on all the nodes in $\mathbf{r}^{(l)}$, whose time complexity is $O(VKA \prod_{l=2}^L D^{(l)})$ per epoch. Both time complexities are $\prod_{l=2}^L D^{(l)}$ times higher than batch GCN, where $\prod_{l=2}^L D^{(l)} = 2$ if we sample 2 neighbors per node and there are 2 GCN layers.

Our algorithm requires an additional $O(VLA)$ space to store historical activations. However, as implemented in our code, the history can be stored in main memory along with the data, which should be larger.

4. Theoretical Results

Besides smaller variance, CV also has stronger theoretical guarantees than NS. In this section, we present two theorems. The first states that if the model parameters are fixed, e.g., during testing, CV produces exact predictions after L epochs; and the second establishes the convergence towards a local optimum regardless of the neighbor sampling size.

In this section, we assume that the algorithm is run by epochs, where each epoch contains I iterations, and in each iteration we want to compute the stochastic gradient w.r.t. nodes in \mathcal{V}_i . We ensure that the activations of all nodes are computed at least one in each epoch, so that the staleness of the history is bounded. We use the subscript i for iteration number and ${}_{CV}$ to distinguish CV from the exact algorithm, i.e., $Z_i^{(l)}$ and $H_i^{(l)}$, W_i , and $g_i(W_i) := \frac{1}{|\mathcal{V}_i|} \sum_{v \in \mathcal{V}_i} \nabla f(y_v, z_{i,v}^{(L)})$ are the activations, model weights, and stochastic gradients obtained by the exact algorithm; and $Z_{CV,i}^{(l)}$, $H_{CV,i}^{(l)}$, and $g_{CV,i}(W_i)$ are their CV counterparts. $\nabla \mathcal{L}(W_i) = \frac{1}{|\mathcal{V}_L|} \sum_{v \in \mathcal{V}_L} \nabla f(y_v, z_v^{(L)})$ is the deterministic batch gradient computed by the exact algorithm. The subscript i may be omitted for the exact algorithm if W_i is a constant sequence. We let $[L] = \{0, \dots, L\}$ and $[L]_+ = \{1, \dots, L\}$.

4.1. Exact Testing

The following theorem reveals the connection between the exact predictions and the approximate predictions by CV. The proof can be found in Appendix B.

Theorem 1. *For a constant sequence of $W_i = W$ and any $i > LI$ (i.e., after L epochs), the activations computed by CV are exact, i.e., $Z_{CV,i}^{(l)} = Z^{(l)}$ for each $l \in [L]$ and $H_{CV,i}^{(l)} = H^{(l)}$ for each $l \in [L-1]$.*

Theorem 1 shows that at testing time, we can run forward propagation with CV for L epochs and get exact prediction. This outperforms NS, which cannot recover the exact prediction unless the neighbor sample size goes to infinity. Comparing with directly making exact predictions by an exact batch algorithm, CV is more scalable because it does not need to load the entire graph into memory.

4.2. Convergence Guarantee

The following theorem shows that SGD training with the approximated gradients $g_{CV,i}(W_i)$ still converges to a local optimum, regardless of the neighbor sampling size $D^{(l)}$.

Theorem 2. *Assume that (1) the activation $\sigma(\cdot)$ is ρ -Lipschitz, (2) the gradient of the cost function $\nabla_z f(y, z)$ is ρ -Lipschitz and bounded, (3) $\|g_{CV,\mathcal{V}}(W)\|_\infty$, $\|g(W)\|_\infty$, and $\|\nabla \mathcal{L}(W)\|_\infty$ are all bounded by $G > 0$ for all \hat{P} , \mathcal{V} and W . (4) The loss $\mathcal{L}(W)$ is ρ -smooth, i.e., $|\mathcal{L}(W_2) - \mathcal{L}(W_1)| -$*

$\langle \nabla \mathcal{L}(W_1), W_2 - W_1 \rangle \leq \frac{\rho}{2} \|W_2 - W_1\|_F^2 \forall W_1, W_2$, where $\langle A, B \rangle = \text{tr}(A^\top B)$ is the inner product of matrix A and matrix B . (5) The loss $\mathcal{L}(W) \geq \mathcal{L}_$ is bounded below. Then, there exists $K > 0$, s.t., $\forall N > LI$, if we run SGD for $R \leq N$ iterations, where R is chosen uniformly from $[N]_+$, we have*

$$\mathbb{E}_R \|\nabla \mathcal{L}(W_R)\|_F^2 \leq 2 \frac{\mathcal{L}(W_1) - \mathcal{L}_* + K + \rho K}{\sqrt{N}},$$

for the updates $W_{i+1} = W_i - \gamma g_{CV,i}(W_i)$ and the step size $\gamma = \min\{\frac{1}{\rho}, \frac{1}{\sqrt{N}}\}$.

Particularly, $\lim_{N \rightarrow \infty} \mathbb{E}_R \|\nabla \mathcal{L}(W_R)\|^2 = 0$. Therefore, our algorithm converges to a local optimum W where the batch gradient $\nabla \mathcal{L}(W) = 0$. The full proof is in Appendix C. For short, we show that $g_{CV,i}(W_i)$ is unbiased as $i \rightarrow \infty$, and then show that SGD with such asymptotically unbiased gradients converges to a local optimum.

Theorem 2 generalizes to graph attention networks (GAT) (Veličković et al., 2018). We leave the variance reduced stochastic estimators for GAT, and discussions on the convergence of GAT and other models in Appendix C.5.

5. Handling Dropout of Features

In this section, we consider introducing a third source of randomness, the random dropout of features (Srivastava et al., 2014), which is adopted in various GCN models as a regularization (Kipf & Welling, 2017; Veličković et al., 2018). With dropout, the GCN layer becomes $Z^{(l+1)} = M \circ (PH^{(l)})W^{(l)}$, where $M_{ij} \sim \text{Bern}(p)$ are i.i.d. Bernoulli random variables, and \circ is the element-wise product. Let \mathbb{E}_M be the expectation over dropout masks.

With dropout, all the activations $h_v^{(l)}$ are random variables whose randomness comes from dropout, even in the exact algorithm Eq. (1). We want to design a cheap estimator for the random variable $(PH^{(l)})_u = \sum_{v \in \mathfrak{n}(u)} P_{uv} h_v^{(l)}$, based on a stochastic neighborhood $\hat{\mathfrak{n}}^{(l)}(u)$. An ideal estimator should have the same distribution with $(PH^{(l)})_u$. However, such an estimator is difficult to design. Instead, we develop an estimator $\text{CVD}_u^{(l)}$ that eventually has the same mean and variance with $(PH^{(l)})_u$, i.e., $\mathbb{E}_{\hat{\mathfrak{n}}^{(l)}(u)} \mathbb{E}_M \text{CVD}_u^{(l)} = \mathbb{E}_M (PH^{(l)})_u$ and $\text{Var}_{\hat{\mathfrak{n}}^{(l)}(u)} \text{Var}_M \text{CVD}_u^{(l)} = \text{Var}_M (PH^{(l)})_u$.

5.1. Control Variate for Dropout

With dropout, $\Delta h_v^{(l)} = h_v^{(l)} - \bar{h}_v^{(l)}$ is not necessarily small even if $\bar{h}_v^{(l)}$ and $h_v^{(l)}$ have the same distribution. We develop another stochastic approximation algorithm, *control variate for dropout* (CVD), that works well with dropout.

Estimator	VNS	VD
Exact	0	$S_u^{(l)}$
NS	$\frac{C_u^{(l)}}{2D^{(l)}} \sum_{v_1, v_2 \in \mathbf{n}(u)} (P_{uv_1} \mu_{v_1}^{(l)} - P_{uv_2} \mu_{v_2}^{(l)})^2$	$\frac{n(u)}{D^{(l)}} S_u^{(l)}$
CV	$\frac{C_u^{(l)}}{2D^{(l)}} \sum_{v_1, v_2 \in \mathbf{n}(u)} (P_{uv_1} \Delta \mu_{v_1}^{(l)} - P_{uv_2} \Delta \mu_{v_2}^{(l)})^2$	$\left(3 + \frac{n(u)}{D^{(l)}}\right) S_u^{(l)}$
CVD	$\frac{C_u^{(l)}}{2D^{(l)}} \sum_{v_1, v_2 \in \mathbf{n}(u)} (P_{uv_1} \Delta \mu_{v_1}^{(l)} - P_{uv_2} \Delta \mu_{v_2}^{(l)})^2$	$S_u^{(l)}$

Table 2. Variance from neighbor sampling (VNS) and variance from dropout (ND) of different estimators.

Our method is based on the weight scaling procedure (Srivastava et al., 2014) to approximately compute the mean $\mu_v^{(l)} := \mathbb{E}_M [h_v^{(l)}]$. That is, along with the dropout model, we can run a copy of the model without dropout to obtain the mean $\mu_v^{(l)}$, as illustrated in Fig. 1(d). We obtain a stochastic approximation by separating the mean and variance

$$\begin{aligned} (PH^{(l)})_u &= \sum_{v \in \mathbf{n}(u)} P_{uv} (\mathring{h}_v^{(l)} + \Delta \mu_v^{(l)} + \bar{\mu}_v^{(l)}) \approx \text{CVD}_u^{(l)} \\ &:= \sqrt{R} \sum_{v \in \hat{\mathbf{n}}} P_{uv} \mathring{h}_v^{(l)} + R \sum_{v \in \hat{\mathbf{n}}} P_{uv} \Delta \mu_v^{(l)} + \sum_{v \in \mathbf{n}(u)} P_{uv} \bar{\mu}_v^{(l)}, \end{aligned}$$

where we define $\mathbf{n} = \hat{\mathbf{n}}^{(l)}(u)$, $R = n(u)/D^{(l)}$ for short, $\mathring{h}_v^{(l)} = h_v^{(l)} - \mu_v^{(l)}$, $\bar{\mu}_v^{(l)}$ is the historical mean activation, obtained by storing $\mu_v^{(l)}$ instead of $h_v^{(l)}$, and $\Delta \mu_v^{(l)} = \mu_v^{(l)} - \bar{\mu}_v^{(l)}$. We separate $h_v^{(l)}$ as three terms, the latter two terms on $\mu_v^{(l)}$ do not have the randomness from dropout, and $\mu_v^{(l)}$ are treated as if $h_v^{(l)}$ for the CV estimator. The first term has zero mean w.r.t. dropout, i.e., $\mathbb{E}_M \mathring{h}_v^{(l)} = 0$. We have $\mathbb{E}_{\hat{\mathbf{n}}^{(l)}(u)} \mathbb{E}_M \text{CVD}_u^{(l)} = 0 + \sum_{v \in \mathbf{n}(u)} P_{uv} (\Delta \mu_v^{(l)} + \bar{\mu}_v^{(l)}) = \mathbb{E}_M (PH^{(l)})_u$, i.e., the estimator is unbiased, and we shall see that the estimator eventually has the correct variance if $h_v^{(l)}$'s are uncorrelated in Sec. 5.2.

5.2. Variance Analysis

We analyze the variance under the assumption that the node activations are uncorrelated, i.e., $\text{Cov}_M [h_{v_1}^{(l)}, h_{v_2}^{(l)}] = 0, \forall v_1 \neq v_2$. We report the correlation between nodes empirically in Appendix G. To facilitate the analysis of the variance, we introduce two propositions proven in Appendix A. The first helps the derivation of the dropout variance; and the second implies that we can treat the variance introduced by neighbor sampling and by dropout separately.

Proposition 2. *If $\hat{\mathbf{n}}^{(l)}(u)$ contains $D^{(l)}$ samples from the set $\mathbf{n}(u)$ without replacement, x_1, \dots, x_V are random variables, $\forall v, \mathbb{E}[x_v] = 0$ and $\forall v_1 \neq v_2, \text{Cov}[x_{v_1}, x_{v_2}] = 0$, then $\text{Var}_{X, \hat{\mathbf{n}}^{(l)}(u)} \left[\frac{n(u)}{D^{(l)}} \sum_{v \in \hat{\mathbf{n}}^{(l)}(u)} x_v \right] = \frac{n(u)}{D^{(l)}} \sum_{v \in \mathbf{n}(u)} \text{Var}[x_v]$.*

Proposition 3. *X and Y are two random variables, and*

$f(X, Y)$ and $g(Y)$ are two functions. If $E_X f(X, Y) = 0$, then $\text{Var}_{X, Y} [f(X, Y) + g(Y)] = \text{Var}_{X, Y} f(X, Y) + \text{Var}_Y g(Y)$.

By Proposition 3, $\text{Var}_{\hat{\mathbf{n}}} \text{Var}_M \text{CVD}_u^{(l)}$ can be written as the sum of $\text{Var}_{\hat{\mathbf{n}}} \text{Var}_M \left[\sqrt{R} \sum_{v \in \hat{\mathbf{n}}} P_{uv} \mathring{h}_v^{(l)} \right]$ and $\text{Var}_{\hat{\mathbf{n}}} \left[R \sum_{v \in \hat{\mathbf{n}}} P_{uv} \Delta \mu_v^{(l)} + \sum_{v \in \mathbf{n}(u)} P_{uv} \bar{\mu}_v^{(l)} \right]$. We refer the first term as the variance from dropout (VD) and the second term as the variance from neighbor sampling (VNS). Ideally, VD should equal to the variance of $(PH^{(l)})_u$ and VNS should be zero. VNS can be derived by replicating the analysis in Sec. 3.2, and replacing h with μ . Let $s_v^{(l)} = \text{Var}_M h_v^{(l)} = \text{Var}_M \mathring{h}_v^{(l)}$, and $S_u^{(l)} = \text{Var}_M (PH^{(l)})_u = \sum_{v \in \mathbf{n}(u)} P_{uv}^2 s_v^{(l)}$. By Proposition 2, VD of $\text{CVD}_u^{(l)}$ is $\sum_{v \in \mathbf{n}(u)} P_{uv}^2 \text{Var} [\mathring{h}_v^{(l)}] = S_u^{(l)}$, which equals with the VD of the exact estimator as desired.

We summarize the estimators and their variances in Table 2, where the derivations are in Appendix A. As in Sec. 3.2, VNS of CV and CVD depends on $\Delta \mu_v$, which converges to zero as the training progresses, while VNS of NS depends on the non-zero μ_v . On the other hand, CVD is the only estimator except the exact one that gives correct VD.

5.3. Preprocessing Strategy

There are two possible models adopting dropout, $Z^{(l+1)} = P(M \circ H^{(l)})W^{(l)}$ or $Z^{(l+1)} = M \circ (PH^{(l)})W^{(l)}$. The difference is whether the dropout layer is before or after neighbor averaging. Kipf & Welling (2017) adopt the former one, and we adopt the latter one, while the two models perform similarly in practice, as we shall see in Sec. 6.1. The advantage of the latter model is that we can preprocess $U^{(0)} = PH^{(0)} = PX$ and takes $U^{(0)}$ as the new input. In this way, the actual number of graph convolution layers is reduced by one — the first layer is merely a fully-connected layer instead of a graph convolution one. Since most GCNs only have two graph convolution layers (Kipf & Welling, 2017; Hamilton et al., 2017a), this gives a significant reduction of the receptive field size and speeds up the computation. We refer this optimization as the preprocessing strategy.

Dataset	M0	M1	M1+PP
Citeseer	70.8 ± .1	70.9 ± .2	70.9 ± .2
Cora	81.7 ± .5	82.0 ± .8	81.9 ± .7
PubMed	79.0 ± .4	78.7 ± .3	78.9 ± .5
NELL	-	64.9 ± 1.7	64.2 ± 4.6
PPI	97.9 ± .04	97.8 ± .05	97.6 ± .09
Reddit	96.2 ± .04	96.3 ± .07	96.3 ± .04

Table 3. Testing accuracy of different algorithms and models after fixed number of epochs. Our implementation does not support M0 on NELL so the result is not reported.

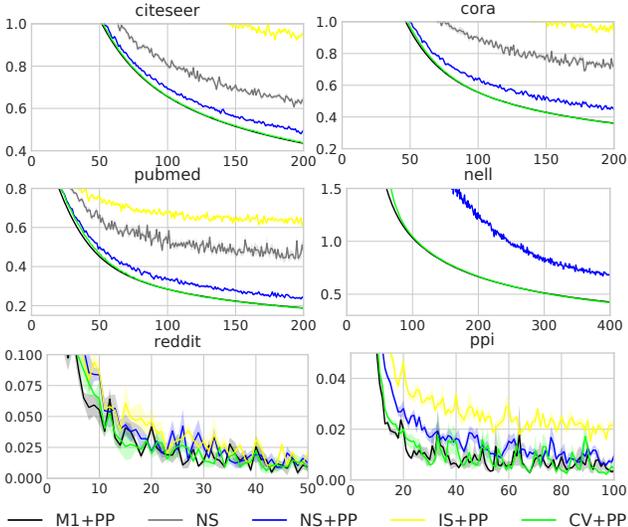


Figure 2. Comparison of training loss with respect to number of epochs without dropout. The CV+PP curve overlaps with the Exact curve in the first four datasets. The training loss of NS and IS+PP are not shown on some datasets because they are too high.

6. Experiments

We examine the variance and convergence of our algorithms empirically on six datasets, including Citeseer, Cora, PubMed and NELL from Kipf & Welling (2017) and Reddit, PPI from Hamilton et al. (2017a), with the same train / validation / test splits, as summarized in Table 1. To measure the predictive performance, we report Micro-F1 for the multi-label PPI dataset, and accuracy for all the other multi-class datasets. The model is GCN for the former 4 datasets and GraphSAGE-mean (Hamilton et al., 2017a) for the latter 2 datasets, see Appendix E for the details on the architectures. We repeat the convergence experiments 10 times on Citeseer, Cora, PubMed and NELL, and 5 times on Reddit and PPI. The experiments are done on a Titan X (Maxwell) GPU.

6.1. Impact of Preprocessing

We first examine the impact of switching the order of dropout and computing neighbor averaging in Sec. 5.3. Let M0 be the $Z^{(l+1)} = P(M \circ H^{(l)})W^{(l)}$ model by (Kipf & Welling, 2017), and M1 be our $Z^{(l+1)} = M \circ (PH^{(l)})W^{(l)}$

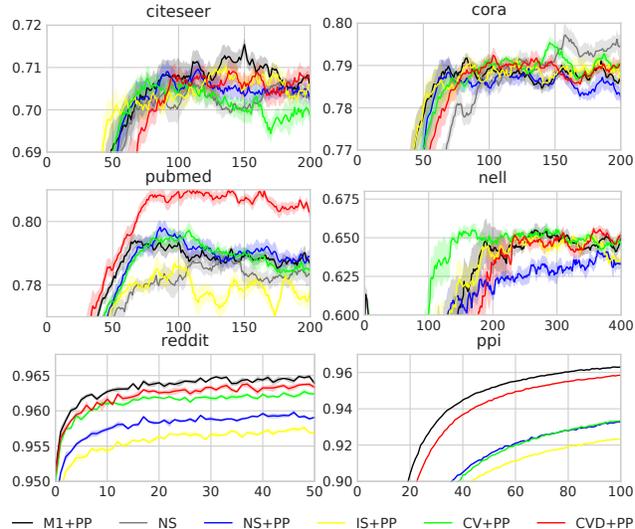


Figure 3. Comparison of validation accuracy with respect to number of epochs. NS converges to 0.94 on the Reddit dataset and 0.6 on the PPI dataset.

model, we compare three settings: M0 and M1 are exact algorithms without any neighbor sampling, and M1+PP samples a large number of $D^{(l)} = 20$ neighbors and preprocesses $PH^{(0)}$ so that the first neighbor averaging is exact. In Table 3 we can see that all the three settings performs similarly, i.e., switching the order does not affect the predictive performance. Therefore, we use the fastest M1+PP as the exact baseline in following convergence experiments.

6.2. Convergence Results

Having the M1+PP algorithm as an exact baseline, the next goal is reducing the time complexity per epoch to make it comparable with the time complexity of MLP, by setting $D^{(l)} = 2$. We cannot set $D^{(l)} = 1$ because GraphSAGE explicitly need the activation of a node itself besides the average of its neighbors. Four approximate algorithms are included for comparison: (1) NS, which adopts the NS estimator with no preprocessing. (2) NS+PP, which is same with NS but uses preprocessing. (3) CV+PP, which adopts the CV estimator and preprocessing. (4) CVD+PP, which uses the CVD estimator. All the four algorithms have similar low time complexity per epoch with $D^{(l)} = 2$, while M1+PP takes $D^{(l)} = 20$. We study how much convergence speed per epoch and model quality do these approximate algorithms sacrifice comparing with the M1+PP baseline.

We set the dropout rate as zero and plot the training loss with respect to number of epochs as Fig. 2. We can see that CV+PP can always reach the same training loss with M1+PP, while NS, NS+PP and IS+PP have higher training losses because of their biased gradients. CVD+PP is not included because it is the same with CV+PP when the dropout rate is zero. The results matches the conclusion of Theorem 2, which states that training with the CV estimator converges

Alg.	Valid. acc.	Epochs	Time (s)
M1+PP	96.0	$4.8 \pm .7$	252 ± 37
NS	$94.4 \pm .01$	100	445 ± 14
NS+PP	96.0	39.8 ± 11	161 ± 47
IS+PP	$95.8 \pm .1$	50	251 ± 6
CV+PP	96.0	7.6 ± 1.6	39 ± 8
CVD+PP	96.0	6.8 ± 1.3	37 ± 7

Table 4. Time complexity comparison of different algorithms on the Reddit dataset.

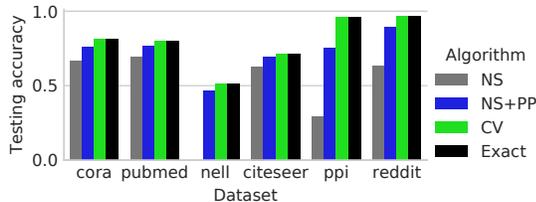


Figure 4. Comparison of the accuracy of different testing algorithms. The y-axis is Micro-F1 for PPI and accuracy otherwise.

to a local optimum of Exact, regardless of $D^{(l)}$.

Next, we turn dropout on and compare the validating accuracy obtained by the model trained with different algorithms at each epoch. Regardless of the training algorithm, the exact algorithm is used for computing predictions on the validating set. The result is shown in Fig. 3. We find that when dropout is present, CVD+PP is the only algorithm that can reach comparable validation accuracy with the exact algorithm on all datasets. Furthermore, its convergence speed with respect to the number of epochs is comparable with M1+PP, implying almost no loss of the convergence speed despite its $D^{(l)}$ is 10 times smaller. This is already the best we can expect - comparable time complexity with MLP, yet similar model quality with GCN. CVD+PP performs much better than M1+PP on the PubMed dataset, we suspect it finds a better local optimum. Meanwhile, the simpler CV+PP also reaches a comparable accuracy with M1+PP for all datasets except PPI. IS+PP works worse than NS+PP on the Reddit and PPI datasets, perhaps because sometimes nodes can have no neighbor selected, as we mentioned in Sec. 2.4. Our accuracy result for IS+PP can match the result reported by Chen et al. (2018), while their NS baseline, GraphSAGE (Hamilton et al., 2017a), does not implement the preprocessing technique in Sec. 5.3.

6.3. Further Analysis on Time Complexity, Testing Accuracy and Variance

Table 4 reports the average number of epochs and time to reach a given 96% validation accuracy on the largest Reddit dataset. Sparse and dense computations are defined in Sec. 3.4. We found that CVD+PP is about 7 times faster than M1+PP due to the significantly reduced receptive field size. Meanwhile, NS and IS+PP does not converge to the given accuracy.

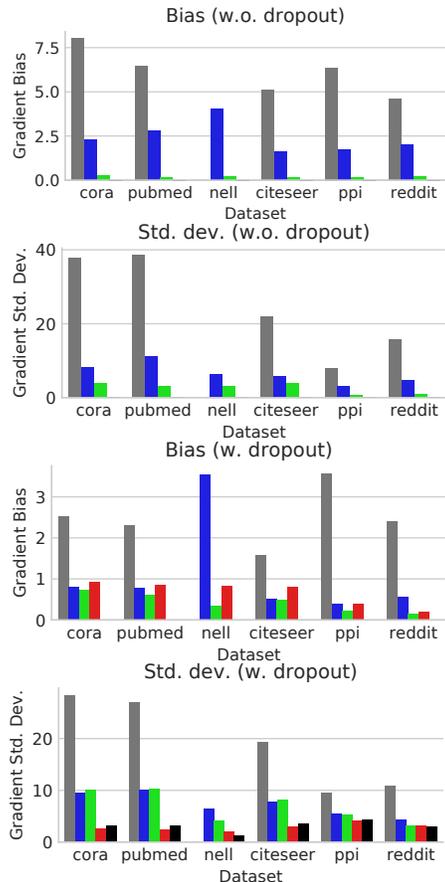


Figure 5. Bias and standard deviation of the gradient for different algorithms during training.

We compare the quality of the predictions made by different algorithms, using the *same* model trained with M1+PP in Fig. 4. As Theorem 1 states, CV reaches the same testing accuracy as the exact algorithm, while NS and NS+PP perform much worse.

Finally, we compare the average bias and variance of the gradients per dimension for first layer weights relative to the magnitude of the weights in Fig. 5. For models without dropout, the gradient of CV+PP is almost unbiased. For models with dropout, the bias and variance of CV+PP and CVD+PP are usually smaller than NS and NS+PP.

7. Conclusions

The large receptive field size of GCN hinders its fast stochastic training. In this paper, we present control variate based algorithms to reduce the receptive field size. Our algorithms can achieve comparable convergence speed with the exact algorithm even the neighbor sampling size $D^{(l)} = 2$, so that the per-epoch cost of training GCN is comparable with training MLPs. We also present strong theoretical guarantees, including exact prediction and the convergence to a local optimum. Our code is released at https://github.com/thu-ml/stochastic_gcnn.

Acknowledgements

We thank Shuyu Cheng for his help in proofreading. This work was supported by NSFC Projects (Nos. 61620106010, 61621136008, 61332007), Beijing NSF Project (No. L172037), Tiangong Institute for Intelligent Computing, NVIDIA NVAIL Program, Siemens and Intel. L.S. was also supported in part by NSF IIS-1218749, NIH BIGDATA 1R01GM108341, NSF CAREER IIS-1350983, NSF IIS-1639792 EAGER, NSF CNS-1704701, ONR N00014-15-1-2340, Intel ISTC, NVIDIA and Amazon AWS.

References

- Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., et al. Tensorflow: A system for large-scale machine learning. In *OSDI*, volume 16, pp. 265–283, 2016.
- Berg, R. v. d., Kipf, T. N., and Welling, M. Graph convolutional matrix completion. *arXiv preprint arXiv:1706.02263*, 2017.
- Chen, J., Ma, T., and Xiao, C. Fastgcn: Fast learning with graph convolutional networks via importance sampling. In *ICLR*, 2018.
- Grover, A. and Leskovec, J. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 855–864. ACM, 2016.
- Hamilton, W., Ying, Z., and Leskovec, J. Inductive representation learning on large graphs. In *Advances in Neural Information Processing Systems*, pp. 1025–1035, 2017a.
- Hamilton, W. L., Ying, R., and Leskovec, J. Representation learning on graphs: Methods and applications. *arXiv preprint arXiv:1709.05584*, 2017b.
- Kipf, T. N. and Welling, M. Variational graph auto-encoders. *arXiv preprint arXiv:1611.07308*, 2016.
- Kipf, T. N. and Welling, M. Semi-supervised classification with graph convolutional networks. In *ICLR*, 2017.
- LeCun, Y., Bengio, Y., et al. Convolutional networks for images, speech, and time series. *The handbook of brain theory and neural networks*, 3361(10):1995, 1995.
- Perozzi, B., Al-Rfou, R., and Skiena, S. Deepwalk: Online learning of social representations. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 701–710. ACM, 2014.
- Ripley, B. D. *Stochastic simulation*, volume 316. John Wiley & Sons, 2009.
- Schlichtkrull, M., Kipf, T. N., Bloem, P., Berg, R. v. d., Titov, I., and Welling, M. Modeling relational data with graph convolutional networks. *arXiv preprint arXiv:1703.06103*, 2017.
- Srivastava, N., Hinton, G. E., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. Dropout: a simple way to prevent neural networks from overfitting. *Journal of machine learning research*, 15(1):1929–1958, 2014.
- Tang, J., Qu, M., Wang, M., Zhang, M., Yan, J., and Mei, Q. Line: Large-scale information network embedding. In *Proceedings of the 24th International Conference on World Wide Web*, pp. 1067–1077. International World Wide Web Conferences Steering Committee, 2015.
- Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., and Bengio, Y. Graph attention networks. In *ICLR*, 2018.

Stochastic Training of Graph Convolutional Networks with Variance Reduction: Supplementary Material

A Derivation of the variance

The following proposition is widely used in this section.

Proposition A. *Let X_1, \dots, X_N are random variables, then*

$$\text{Var} \left[\sum_{i=1}^N X_i \right] = \sum_{i=1}^N \sum_{j=1}^N \text{Cov} [X_i, X_j].$$

Proof.

$$\begin{aligned} \text{Var} \left[\sum_{i=1}^N X_i \right] &= \mathbb{E} \left[\sum_{i=1}^N \sum_{j=1}^N X_i X_j \right] - \left(\mathbb{E} \sum_{i=1}^N X_i \right)^2 \\ &= \mathbb{E} \left[\sum_{i=1}^N \sum_{j=1}^N \left(X_i X_j - \frac{1}{N} \mathbb{E} \sum_{i=1}^N X_i \right) \right] \\ &= \sum_{i=1}^N \sum_{j=1}^N \text{Cov} [X_i, X_j]. \end{aligned}$$

□

We begin with the proof for the three propositions in the main text.

Proposition 1. *If $\hat{\mathbf{n}}^{(l)}(u)$ contains $D^{(l)}$ samples from $\mathbf{n}(u)$ without replacement, then $\text{Var}_{\hat{\mathbf{n}}^{(l)}(u)} \left[\frac{n(u)}{D^{(l)}} \sum_{v \in \hat{\mathbf{n}}^{(l)}(u)} x_v \right] = \frac{C_u^{(l)}}{2D^{(l)}} \sum_{v_1 \in \mathbf{n}(u)} \sum_{v_2 \in \mathbf{n}(u)} (x_{v_1} - x_{v_2})^2$, where $C_u^{(l)} = 1 - (D^{(l)} - 1)/(n(u) - 1)$.*

Proof. We denote the $D^{(l)}$ samples in the set as $v_1, \dots, v_{D^{(l)}}$. Let $\bar{x} = \frac{1}{n(u)} \sum_{v \in \mathbf{n}(u)} x_v$,

then

$$\begin{aligned}
& \text{Var}_{\hat{\mathbf{n}}^{(l)}(u)} \left[\frac{n(u)}{D^{(l)}} \sum_{v \in \hat{\mathbf{n}}^{(l)}(u)} x_v \right] \\
&= \text{Var}_{v_1, \dots, v_{D^{(l)}}} \left[\frac{n(u)}{D^{(l)}} \sum_{i=1}^{D^{(l)}} x_{v_i} \right] \\
&= \left(\frac{n(u)}{D^{(l)}} \right)^2 \sum_{i=1}^{D^{(l)}} \sum_{j=1}^{D^{(l)}} \text{Cov}_{v_1, \dots, v_{D^{(l)}}} [x_{v_i}, x_{v_j}] \\
&= \left(\frac{n(u)}{D^{(l)}} \right)^2 \left\{ \sum_{i=1}^{D^{(l)}} \text{Var}_{v_i} [x_{v_i}^2] + \sum_{i \neq j} \text{Cov}_{v_i, v_j} [x_{v_i}, x_{v_j}] \right\} \\
&= \left(\frac{n(u)}{D^{(l)}} \right)^2 \left\{ \frac{D^{(l)}}{n(u)} \sum_{v \in \mathbf{n}(u)} (x_v - \bar{x})^2 + \frac{D^{(l)}(D^{(l)} - 1)}{n(u)(n(u) - 1)} \left[\sum_{i, j \in \mathbf{n}(u)} (x_i - \bar{x})(x_j - \bar{x}) - \sum_{i \in \mathbf{n}(u)} (x_i - \bar{x})^2 \right] \right\} \\
&= \frac{n(u)}{D^{(l)}} \left(1 - \frac{D^{(l)} - 1}{n(u) - 1} \right) \left(\sum_{v \in \mathbf{n}(u)} x_v^2 - n(u) \bar{x}^2 \right) \\
&= \frac{1}{2D^{(l)}} \left(1 - \frac{D^{(l)} - 1}{n(u) - 1} \right) \left(2n(u) \sum_{v \in \mathbf{n}(u)} x_v^2 - \sum_{v_1, v_2 \in \mathbf{n}(u)} 2x_{v_1} x_{v_2} \right) \\
&= \frac{C_u^{(l)}}{2D^{(l)}} \sum_{v_1, v_2 \in \mathbf{n}(u)} (x_{v_1} - x_{v_2})^2.
\end{aligned}$$

□

Proposition 2. *If $\hat{\mathbf{n}}^{(l)}(u)$ contains $D^{(l)}$ samples from the set $\mathbf{n}(u)$ without replacement, x_1, \dots, x_V are random variables, $\forall v, \mathbb{E}[x_v] = 0$ and $\forall v_1 \neq v_2, \text{Cov}[x_{v_1}, x_{v_2}] = 0$, then $\text{Var}_{X, \hat{\mathbf{n}}^{(l)}(u)} \left[\frac{n(u)}{D^{(l)}} \sum_{v \in \hat{\mathbf{n}}^{(l)}(u)} x_v \right] = \frac{n(u)}{D^{(l)}} \sum_{v \in \mathbf{n}(u)} \text{Var}[x_v]$.*

Proof.

$$\begin{aligned}
& \text{Var}_{X, \hat{\mathbf{n}}^{(l)}(u)} \left[\frac{n(u)}{D^{(l)}} \sum_{v \in \hat{\mathbf{n}}^{(l)}(u)} x_v \right] \\
&= \left(\frac{n(u)}{D^{(l)}} \right)^2 \mathbb{E}_X \left\{ \sum_{i=1}^{D^{(l)}} \mathbb{E}_{\hat{\mathbf{n}}^{(l)}(u)} x_{v_i}^2 + \sum_{i \neq j} \mathbb{E}_{\hat{\mathbf{n}}^{(l)}(u)} x_{v_i} x_{v_j} \right\} \\
&= \left(\frac{n(u)}{D^{(l)}} \right)^2 \mathbb{E}_X \left\{ \frac{D^{(l)}}{n(u)} \sum_{i \in \mathbf{n}(u)} x_i^2 + \frac{D^{(l)}(D^{(l)} - 1)}{n(u)(n(u) - 1)} \sum_{i, j \in \mathbf{n}(u), i \neq j} x_i x_j \right\} \\
&= \frac{n(u)}{D^{(l)}} \sum_{i \in \mathbf{n}(u)} \text{Var} [x_i].
\end{aligned}$$

□

Proposition 3. *X and Y are two random variables, and f(X, Y) and g(Y) are two functions. If $E_X f(X, Y) = 0$, then $\text{Var}_{X, Y} [f(X, Y) + g(Y)] = \text{Var}_{X, Y} f(X, Y) + \text{Var}_Y g(Y)$.*

Proof.

$$\text{Var}_{X, Y} [f(X, Y) + g(Y)] = \text{Var}_{X, Y} f(X, Y) + \text{Var}_Y g(Y) + 2\text{Cov}_{X, Y} [f(X, Y), g(Y)],$$

where

$$\begin{aligned}
\text{Cov}_{X, Y} [f(X, Y), g(Y)] &= \mathbb{E}_Y \mathbb{E}_X [(f(X, Y) - \mathbb{E}_{X, Y} f(X, Y))(g(Y) - E_Y g(Y))] \\
&= \mathbb{E}_Y [(\mathbb{E}_X f(X, Y) - 0)(g(Y) - E_Y g(Y))] \\
&= \mathbb{E}_Y [0(g(Y) - E_Y g(Y))] = 0.
\end{aligned}$$

□

Then, we derive the variance of the estimators with dropout is present.

A.1 Variance of the exact estimator

$$\text{Var}_M \left[\sum_{v \in \mathbf{n}(u)} P_{uv} h_v^{(l)} \right] = \text{Var}_M \left[\sum_{v \in \mathbf{n}(u)} P_{uv} \mathring{h}_v^{(l)} \right] = \sum_{v \in \mathbf{n}(u)} P_{uv}^2 \text{Var}_M \left[\mathring{h}_v^{(l)} \right] = S_u^{(l)}.$$

A.2 Variance of the NS estimator

$$\begin{aligned}
& \text{Var}_{\hat{\mathbf{n}}^{(l)}(u),M} \left[\text{NS}_u^{(l)} \right] \\
&= \text{Var}_{\hat{\mathbf{n}}^{(l)}(u),M} \left[\frac{n(u)}{D^{(l)}} \sum_{v \in \hat{\mathbf{n}}^{(l)}(u)} P_{uv} h_v^{(l)} \right] \\
&= \text{Var}_{\hat{\mathbf{n}}^{(l)}(u),M} \left[\frac{n(u)}{D^{(l)}} \sum_{v \in \hat{\mathbf{n}}^{(l)}(u)} P_{uv} (\dot{h}_v^{(l)} + \mu_v^{(l)}) \right] \\
&= \text{Var}_{\hat{\mathbf{n}}^{(l)}(u),M} \left[\frac{n(u)}{D^{(l)}} \sum_{v \in \hat{\mathbf{n}}^{(l)}(u)} P_{uv} \dot{h}_v^{(l)} \right] + \text{Var}_{\hat{\mathbf{n}}^{(l)}(u)} \left[\frac{n(u)}{D^{(l)}} \sum_{v \in \hat{\mathbf{n}}^{(l)}(u)} P_{uv} \mu_v^{(l)} \right],
\end{aligned}$$

where the last equality is by Proposition 3. By Proposition 2, VD is

$$\begin{aligned}
& \text{Var}_{\hat{\mathbf{n}}^{(l)}(u),M} \left[\frac{n(u)}{D^{(l)}} \sum_{v \in \hat{\mathbf{n}}^{(l)}(u)} P_{uv} \dot{h}_v^{(l)} \right] \\
&= \text{Var}_M \text{Var}_{\hat{\mathbf{n}}^{(l)}(u)} \left[\frac{n(u)}{D^{(l)}} \sum_{v \in \hat{\mathbf{n}}^{(l)}(u)} P_{uv} \dot{h}_v^{(l)} \right] \\
&= \frac{n(u)}{D^{(l)}} S_u^{(l)},
\end{aligned}$$

where $S_u^{(l)} = \sum_{v \in \mathbf{n}(u)} \text{Var}_M \left[P_{uv} h_v^{(l)} \right]$ is defined in Sec. 5.2. By Proposition 1, VNS is

$$\text{Var}_{\hat{\mathbf{n}}^{(l)}(u)} \left[\frac{n(u)}{D^{(l)}} \sum_{v \in \hat{\mathbf{n}}^{(l)}(u)} P_{uv} \mu_v^{(l)} \right] = \frac{C_u^{(l)}}{2D^{(l)}} \sum_{v_1, v_2 \in \mathbf{n}(u)} (P_{uv_1} \mu_{v_1}^{(l)} - P_{uv_2} \mu_{v_2}^{(l)})^2.$$

A.3 Variance of the CVD estimator

$$\begin{aligned}
& \text{Var}_{\hat{\mathbf{n}}^{(l)}(u),M} \left[\text{CVD}_u^{(l)} \right] \\
&= \text{Var}_{\hat{\mathbf{n}}^{(l)}(u),M} \left[\sqrt{\frac{n(u)}{D^{(l)}}} \sum_{v \in \hat{\mathbf{n}}^{(l)}(u)} P_{uv} \dot{h}_v^{(l)} + \frac{n(u)}{D^{(l)}} \sum_{v \in \hat{\mathbf{n}}^{(l)}(u)} P_{uv} \Delta \mu_v^{(l)} + \sum_{v \in \mathbf{n}(u)} P_{uv} \bar{\mu}_v^{(l)} \right] \\
&= \text{Var}_{\hat{\mathbf{n}}^{(l)}(u),M} \left[\sqrt{\frac{n(u)}{D^{(l)}}} \sum_{v \in \hat{\mathbf{n}}^{(l)}(u)} P_{uv} \dot{h}_v^{(l)} \right] + \text{Var}_{\hat{\mathbf{n}}^{(l)}(u)} \left[\frac{n(u)}{D^{(l)}} \sum_{v \in \hat{\mathbf{n}}^{(l)}(u)} P_{uv} \Delta \mu_v^{(l)} + \sum_{v \in \mathbf{n}(u)} P_{uv} \bar{\mu}_v^{(l)} \right],
\end{aligned}$$

where the last equality is by Proposition 3. By Proposition 2, VD is

$$\begin{aligned}
& \text{Var}_{\hat{\mathbf{n}}^{(l)}(u), M} \left[\sqrt{\frac{n(u)}{D^{(l)}}} \sum_{v \in \hat{\mathbf{n}}^{(l)}(u)} P_{uv} \dot{h}_v^{(l)} \right] \\
&= \frac{D^{(l)}}{n(u)} \text{Var}_M \text{Var}_{\hat{\mathbf{n}}^{(l)}(u)} \left[\frac{n(u)}{D^{(l)}} \sum_{v \in \hat{\mathbf{n}}^{(l)}(u)} P_{uv} \dot{h}_v^{(l)} \right] \\
&= S_u^{(l)}.
\end{aligned}$$

By Proposition 1, VNS is

$$\begin{aligned}
& \text{Var}_{\hat{\mathbf{n}}^{(l)}(u)} \left[\frac{n(u)}{D^{(l)}} \sum_{v \in \hat{\mathbf{n}}^{(l)}(u)} P_{uv} \Delta \mu_v^{(l)} + \sum_{v \in \mathbf{n}(u)} P_{uv} \bar{\mu}_v^{(l)} \right] \\
&= \text{Var}_{\hat{\mathbf{n}}^{(l)}(u)} \left[\frac{n(u)}{D^{(l)}} \sum_{v \in \hat{\mathbf{n}}^{(l)}(u)} P_{uv} \Delta \mu_v^{(l)} \right] \\
&= \frac{C_u^{(l)}}{2D^{(l)}} \sum_{v_1, v_2 \in \mathbf{n}(u)} (P_{uv_1} \Delta \mu_{v_1}^{(l)} - P_{uv_2} \Delta \mu_{v_2}^{(l)})^2.
\end{aligned}$$

A.4 Variance of the CV estimator

$$\begin{aligned}
& \text{Var}_{\hat{\mathbf{n}}^{(l)}(u), M} \left[\text{CV}_u^{(l)} \right] \\
&= \text{Var}_{\hat{\mathbf{n}}^{(l)}(u), M} \left[\frac{n(u)}{D^{(l)}} \sum_{v \in \hat{\mathbf{n}}^{(l)}(u)} P_{uv} \Delta h_v^{(l)} + \sum_{v \in \mathbf{n}(u)} P_{uv} \bar{h}_v^{(l)} \right] \\
&= \text{Var}_{\hat{\mathbf{n}}^{(l)}(u), M} \left[\frac{n(u)}{D^{(l)}} \sum_{v \in \hat{\mathbf{n}}^{(l)}(u)} P_{uv} \Delta \dot{h}_v^{(l)} + \sum_{v \in \mathbf{n}(u)} P_{uv} \bar{\dot{h}}_v^{(l)} + \frac{n(u)}{D^{(l)}} \sum_{v \in \hat{\mathbf{n}}^{(l)}(u)} P_{uv} \Delta \mu_v^{(l)} + \sum_{v \in \mathbf{n}(u)} P_{uv} \bar{\mu}_v^{(l)} \right] \\
&= \text{Var}_{\hat{\mathbf{n}}^{(l)}(u), M} \left[\frac{n(u)}{D^{(l)}} \sum_{v \in \hat{\mathbf{n}}^{(l)}(u)} P_{uv} \Delta \dot{h}_v^{(l)} + \sum_{v \in \mathbf{n}(u)} P_{uv} \bar{\dot{h}}_v^{(l)} \right] \\
&+ \text{Var}_{\hat{\mathbf{n}}^{(l)}(u)} \left[\frac{n(u)}{D^{(l)}} \sum_{v \in \hat{\mathbf{n}}^{(l)}(u)} P_{uv} \Delta \mu_v^{(l)} + \sum_{v \in \mathbf{n}(u)} P_{uv} \bar{\mu}_v^{(l)} \right],
\end{aligned}$$

where $\Delta \dot{h}_v^{(l)} = (h_v^{(l)} - \mu_v^{(l)}) - (\bar{h}_v^{(l)} - \bar{\mu}_v^{(l)})$, and the last equality is by Proposition 3. The VNS term is the same with CVD's VNS term.

To analyze the VD, we further assume $\bar{\dot{h}}_v^{(l)}$ and $\dot{h}_v^{(l)}$ are i.i.d., so $\mathbb{E}_M \Delta \dot{h}_v^{(l)} = 0$,

$$\mathbb{E}_M(\Delta \dot{h}_v^{(l)})^2 = 2\mathbb{E}_M(\dot{h}_v^{(l)})^2, \text{ and } \mathbb{E}_M \dot{h}_v^{(l)} \Delta \dot{h}_v^{(l)} = \mathbb{E}_M(\dot{h}_v^{(l)})^2.$$

$$\begin{aligned} & \text{Var}_{\hat{\mathbf{n}}^{(l)}(u), M} \left[\frac{n(u)}{D^{(l)}} \sum_{v \in \hat{\mathbf{n}}^{(l)}(u)} P_{uv} \Delta \dot{h}_v^{(l)} + \sum_{v \in \mathbf{n}(u)} P_{uv} \bar{\bar{h}}_v^{(l)} \right] \\ = & \mathbb{E}_{\hat{\mathbf{n}}^{(l)}(u), M} \left\{ \left(\frac{n(u)}{D^{(l)}} \right)^2 \sum_{i, j \in \hat{\mathbf{n}}^{(l)}(u)} P_{ui} P_{uj} \Delta \dot{h}_i^{(l)} \Delta \dot{h}_j^{(l)} + \sum_{i, j \in \mathbf{n}(u)} P_{ui} P_{uj} \bar{\bar{h}}_i^{(l)} \bar{\bar{h}}_j^{(l)} \right. \\ & \left. + 2 \frac{n(u)}{D^{(l)}} \sum_{i \in \hat{\mathbf{n}}^{(l)}(u), j \in \mathbf{n}(u)} P_{ui} P_{uj} \Delta \dot{h}_i^{(l)} \bar{\bar{h}}_j^{(l)} \right\} \\ = & \left(\frac{n(u)}{D^{(l)}} \right)^2 \sum_{i \in \hat{\mathbf{n}}^{(l)}(u)} \mathbb{E}_{\hat{\mathbf{n}}^{(l)}(u), M} (P_{ui} \Delta \dot{h}_i^{(l)})^2 + \sum_{i \in \mathbf{n}(u)} P_{ui}^2 \mathbb{E}_M \bar{\bar{h}}_i^{(l)} + 2 \sum_{ij \in \mathbf{n}(u)} P_{ui} P_{uj} \Delta \dot{h}_i^{(l)} \bar{\bar{h}}_j^{(l)} \\ = & \frac{n(u)}{D^{(l)}} \sum_{i \in \mathbf{n}(u)} P_{ui}^2 \mathbb{E}_M (\dot{h}_i^{(l)})^2 + \sum_{i \in \mathbf{n}(u)} P_{ui}^2 \mathbb{E}_M \bar{\bar{h}}_i^{(l)} + 2 \sum_{i \in \mathbf{n}(u)} P_{ui}^2 \mathbb{E}_M \bar{\bar{h}}_i^{(l)} \\ = & \left(3 + \frac{n(u)}{D^{(l)}} \right) S_u^{(l)}. \end{aligned}$$

B Proof of Theorem 1

Theorem 1. For a constant sequence of $W_i = W$ and any $i > LI$ (i.e., after L epochs), the activations computed by CV are exact, i.e., $Z_{CV,i}^{(l)} = Z^{(l)}$ for each $l \in [L]$ and $H_{CV,i}^{(l)} = H^{(l)}$ for each $l \in [L-1]$.

Proof. We prove by induction. After the first epoch the activation $h_{i,v}^{(0)}$ is at least computed once for each node v . So $\bar{H}_{CV,i}^{(0)} = H_{CV,i}^{(0)} = H^{(0)} = X$ for all $i > I$. Assume that we have $\bar{H}_{CV,i}^{(l)} = H_{CV,i}^{(l)} = H^{(l)}$ for all $i > (l+1)I$. Then for all $i > (l+1)I$

$$Z_{CV,i}^{(l+1)} = \left(\hat{P}_i^{(l)} (H_{CV,i}^{(l)} - \bar{H}_{CV,i}^{(l)}) + P \bar{H}_{CV,i}^{(l)} \right) W^{(l)} = P \bar{H}_{CV,i}^{(l)} W^{(l)} = P H^{(l)} W^{(l)} = Z^{(l+1)}. \quad (1)$$

$$H_{CV,i}^{(l+1)} = \sigma(Z_{CV,i}^{(l+1)}) = H^{(l+1)}$$

After one more epoch, all the activations $h_{CV,i,v}^{(l+1)}$ are computed at least once for each v , so $\bar{H}_{CV,i}^{(l+1)} = H_{CV,i}^{(l+1)} = H^{(l+1)}$ for all $i > (l+2)I$. By induction, we know that after LI steps, we have $\bar{H}_{CV,i}^{(L-1)} = H_{CV,i}^{(L-1)} = H^{(L-1)}$. By Eq. 1 we also have $\bar{Z}_{CV,i}^{(L)} = Z^{(L)}$. \square

C Proof of Theorem 2

We proof Theorem 2 in 3 steps:

1. Lemma 1: For a sequence of weights $W^{(1)}, \dots, W^{(N)}$ which are close to each other, CV's approximate activations are close to the exact activations.
2. Lemma 2: For a sequence of weights $W^{(1)}, \dots, W^{(N)}$ which are close to each other, CV's gradients are close to be unbiased.
3. Theorem 2: An SGD algorithm generates the weights that changes slow enough for the gradient bias goes to zero, so the algorithm converges.

The following proposition is needed in our proof

Proposition B. *Let $\|A\|_\infty = \max_{ij} |A_{ij}|$, then*

- $\|AB\|_\infty \leq \text{col}(A) \|A\|_\infty \|B\|_\infty$, where $\text{col}(A)$ is the number of columns of the matrix A .
- $\|A \circ B\|_\infty \leq \|A\|_\infty \|B\|_\infty$, where \circ is the element wise product.
- $\|A + B\|_\infty \leq \|A\|_\infty + \|B\|_\infty$.

Proof.

$$\begin{aligned} \|AB\|_\infty &= \max_{ij} \left| \sum_k A_{ik} B_{kj} \right| \leq \max_{ij} \left| \sum_k \|A\|_\infty \|B\|_\infty \right| = \text{col}(A) \|A\|_\infty \|B\|_\infty. \\ \|A \circ B\|_\infty &= \max_{ij} |A_{ij} B_{ij}| \leq \max_{ij} \|A\|_\infty \|B\|_\infty = \|A\|_\infty \|B\|_\infty. \\ \|A + B\|_\infty &= \max_{ij} |A_{ij} + B_{ij}| \leq \max_{ij} \{|A_{ij}| + |B_{ij}|\} \leq \max_{ij} |A_{ij}| + \max_{ij} |B_{ij}| = \|A\|_\infty + \|B\|_\infty. \end{aligned}$$

□

We define $C := \max\{\text{col}(P), \text{col}(H^{(0)}), \dots, \text{col}(H^{(L)})\}$ to be the maximum number of columns we can possibly encounter in the proof.

C.1 Single layer GCN

The following proposition states that if the inputs and the weights of an one-layer GCN with CV estimator does not change too much, then its output does not change too much, and is close to the output of an exact one-layer GCN.

Proposition C. *If the activation $\sigma(\cdot)$ is ρ -Lipschitz, for any series of T inputs, weights, and stochastic propagation matrices $(X_i, X_{CV,i}, W_i, \hat{P}_i)_{i=1}^T$, s.t.,*

1. *all the matrices are bound by B , i.e., $\|X_{CV,i}\|_\infty \leq B$, $\|X_i\|_\infty \leq B$, $\|W_i\|_\infty \leq B$ and $\|\hat{P}_i\|_\infty \leq B$,*

2. the differences are bound by ϵ , i.e., $\|X_{CV,i} - X_{CV,j}\|_\infty < \epsilon$, $\|X_{CV,i} - X_i\|_\infty < \epsilon$ and $\|W_i - W_j\|_\infty < \epsilon$,

let $P = \mathbb{E}\hat{P}_i$. If at time i we feed $(X_{CV,i}, W_i, \hat{P}_i)$ to an one-layer GCN with CV estimator to evaluate the prediction for nodes in the minibatch \mathcal{V}_i ,¹

$$Z_{CV,i} = \left(\hat{P}_i(X_{CV,i} - \bar{X}_{CV,i}) + P\bar{X}_{CV,i} \right) W_i, \quad H_{CV,i} = \sigma(Z_{CV,i}).$$

where $\bar{X}_{CV,i}$ is the maintained history at time i , and (X_i, W_i, P) to an one-layer GCN with exact estimator

$$Z_i = PX_iW_i, \quad H_i = \sigma(Z_i),$$

then there exists K that depends on C , B and ρ , s.t. for all $I < i, j \leq T$, where I is the number of iterations per epoch:

1. The outputs does not change too fast: $\|Z_{CV,i} - Z_{CV,j}\|_\infty < K\epsilon$ and $\|H_{CV,i} - H_{CV,j}\|_\infty < K\epsilon$,
2. The outputs are close to the exact output: $\|Z_{CV,i} - Z_i\|_\infty < K\epsilon$ and $\|H_{CV,i} - H_i\|_\infty < K\epsilon$.

Proof. Because for all $i > I$ (i.e., after one epoch), the elements of $\bar{X}_{CV,i}$ are all taken from previous iterations, i.e., $X_{CV,1}, \dots, X_{CV,i-1}$, we know that

$$\|\bar{X}_{CV,i} - X_{CV,i}\|_\infty \leq \max_{j \leq i} \|X_{CV,j} - X_{CV,i}\|_\infty \leq \epsilon \quad (\forall i > I). \quad (2)$$

By triangular inequality, we also know

$$\|\bar{X}_{CV,i} - \bar{X}_{CV,j}\|_\infty < 3\epsilon \quad (\forall i, j > I). \quad (3)$$

$$\|\bar{X}_{CV,i} - X_i\|_\infty < 2\epsilon \quad (\forall i > I). \quad (4)$$

Since $\|X_{CV,1}\|_\infty, \dots, \|X_{CV,T}\|_\infty$ are bounded by B , $\|\bar{X}_{CV,i}\|_\infty$ is also bounded

¹Conceptually we feed the data for all the nodes in \mathcal{V} , but since we only require the predictions for the nodes in \mathcal{V}_i , the algorithm will only fetch the input of a subset of nodes $\subset \mathcal{V}$, and update history for those nodes.

by B for $i > I$. Then,

$$\begin{aligned}
& \|Z_{CV,i} - Z_{CV,j}\|_\infty \\
&= \left\| \left(\hat{P}_i(X_{CV,i} - \bar{X}_{CV,i}) + P\bar{X}_{CV,i} \right) W_i - \left(\hat{P}_j(X_{CV,j} - \bar{X}_{CV,j}) + P\bar{X}_{CV,j} \right) W_j \right\|_\infty \\
&\leq \left\| \hat{P}_i(X_{CV,i} - \bar{X}_{CV,i})W_i - \hat{P}_j(X_{CV,j} - \bar{X}_{CV,j})W_j \right\|_\infty + \rho \|P\bar{X}_{CV,i}W_i - P\bar{X}_{CV,j}W_j\|_\infty \\
&\leq C^2 \left[\left\| \hat{P}_i - \hat{P}_j \right\|_\infty \|X_{CV,i} - \bar{X}_{CV,i}\|_\infty \|W_i\|_\infty \right. \\
&\quad + \left\| \hat{P}_j \right\|_\infty \|X_{CV,i} - \bar{X}_{CV,i} - X_{CV,j} + \bar{X}_{CV,j}\|_\infty \|W_i\|_\infty \\
&\quad + \left\| \hat{P}_j \right\|_\infty \|X_{CV,j} - \bar{X}_{CV,j}\|_\infty \|W_i - W_j\|_\infty \\
&\quad + \|P\|_\infty \|\bar{X}_{CV,i} - \bar{X}_{CV,j}\|_\infty \|W_i\|_\infty \\
&\quad \left. + \|P\|_\infty \|\bar{X}_{CV,j}\|_\infty \|W_i - W_j\|_\infty \right] \\
&\leq C^2 \epsilon \left[\left\| \hat{P}_i - \hat{P}_j \right\|_\infty \|W_i\|_\infty + 2 \left\| \hat{P}_j \right\|_\infty \|W_i\|_\infty + \left\| \hat{P}_j \right\|_\infty \|W_i - W_j\|_\infty \right. \\
&\quad \left. + 3 \left\| \hat{P}_j \right\|_\infty \|W_i\|_\infty + \left\| \hat{P}_j \right\|_\infty \|\bar{X}_{CV,j}\|_\infty \right] \\
&\leq \epsilon C^2 [2B^2 + 2B^2 + 2B^2 + 3B^2 + B^2] \\
&= K_1 \epsilon,
\end{aligned}$$

where $K_1 = 10C^2B^2$, and

$$\begin{aligned}
\|Z_{CV,i} - Z_i\|_\infty &\leq \left\| \left(\hat{P}_i(X_{CV,i} - \bar{X}_{CV,i}) + P(\bar{X}_{CV,i} - X_i) \right) \right\|_\infty \|W_i\|_\infty \\
&\leq C \left(\left\| \hat{P}_i \right\|_\infty \epsilon + 2 \|P\|_\infty \epsilon \right) \|W_i\|_\infty \\
&\leq 3CB^2 \epsilon \\
&= K_2 \epsilon,
\end{aligned}$$

where $K_2 = 3CB^2$. By Lipschitz continuity

$$\begin{aligned}
\|H_{CV,i} - H_{CV,j}\|_\infty &\leq \rho K_1 \epsilon, \\
\|H_{CV,i} - H_i\|_\infty &\leq \rho K_2 \epsilon.
\end{aligned}$$

We just let $K = \max\{\rho K_1, \rho K_2, K_1, K_2\}$. \square

C.2 Lemma 1: Activation of Multi-layer GCN

The following lemma bounds the approximation error of activations in a multi-layer GCN with CV. Intuitively, there is a sequence of slow-changing model parameters (W_i) , where W_i is the model at the i -th iteration. At each iteration i we use GCN with CV and GCN with Exact estimator to compute the activations for the minibatch \mathcal{V}_i , and update the corresponding history. Then after L epochs, the error of the predictions by the CV estimator is bounded by the rate of change of (W_i) , regardless of the stochastic propagation matrix \hat{P}_i .

Lemma 1. Assume all the activations are ρ -Lipschitz, given a fixed dataset X and a sequence of T model weights and stochastic propagation matrices $(W_i, \hat{P}_i)_{i=1}^T$, s.t.,

1. $\|W_i\|_\infty \leq B$ and $\|\hat{P}_i\|_\infty \leq B$,
2. $\|W_i - W_j\|_\infty < \epsilon, \forall i, j$,

let $P = \mathbb{E}\hat{P}_i$. If at time i we feed (X, W_i, \hat{P}_i) to a GCN with CV estimator to evaluate the prediction for nodes in the minibatch \mathcal{V}_i ,

$$Z_{CV,i}^{(l+1)} = \left(\hat{P}_i^{(l)} (H_{CV,i}^{(l)} - \bar{H}_{CV,i}^{(l)}) + P \bar{H}_{CV,i}^{(l)} \right) W_i^{(l)}, \quad H_{CV,i}^{(l+1)} = \sigma(Z_{CV,i}^{(l+1)}).$$

where $\bar{H}_{CV,i}^{(l)}$ is the maintained history at time i , and (X, W_i, P) to a GCN with exact estimator

$$Z_i^{(l+1)} = P H_i^{(l)} W_i^{(l)}, \quad H_i^{(l+1)} = \sigma(Z_i^{(l+1)}),$$

then there exists K that depends on C, B and ρ s.t.,

- $\|H_i^{(L)} - H_{CV,i}^{(L)}\|_\infty < K\epsilon, \forall i > LI, l = 1, \dots, L-1$,
- $\|Z_i^{(L)} - Z_{CV,i}^{(L)}\|_\infty < K\epsilon, \forall i > LI, l = 1, \dots, L$.

Proof. By Proposition C.1, we know there exists $K^{(1)}$, s.t., $\|H_i^{(1)} - H_{CV,i}^{(1)}\| < K^{(1)}\epsilon$ and $\|H_{CV,i}^{(1)} - H_{CV,j}^{(1)}\| < K^{(1)}\epsilon, \forall i > I$.

Repeat this for $L-1$ times, we know there exist $K^{(1)}, \dots, K^{(L)}$, s.t., $\|H_i^{(L)} - H_{CV,i}^{(L)}\| < K\epsilon, \|H_{CV,i}^{(L)} - H_{CV,j}^{(L)}\| < K\epsilon, \|Z_i^{(L)} - Z_{CV,i}^{(L)}\| < K\epsilon$ and $\|Z_{CV,i}^{(L)} - Z_{CV,j}^{(L)}\| < K\epsilon, \forall i > LI$, where $K = \prod_{l=1}^L K^{(l)}$. □

C.3 Lemma 2: Gradient of Multi-layer GCN

We reiterate some notations defined in Sec. 4 of the main text. \mathcal{V}_i is the minibatch of nodes at iteration i that we would like to evaluate the predictions and gradients on. $g_{CV,v}(W_i) := \nabla f(y_v, z_{CV,i,v}^{(L)})$ is the stochastic gradient propagated through the node v by the CV estimator, and $g_{CV,i}(W_i) := \frac{1}{|\mathcal{V}_i|} \sum_{v \in \mathcal{V}_i} \nabla f(y_v, z_{CV,i,v}^{(L)})$ is the minibatch gradient by CV. $g_v(W_i) := \nabla f(y_v, z_v^{(L)})$ is stochastic gradient propagated through the node v by the Exact estimator, and $g_i(W_i) := \frac{1}{|\mathcal{V}_i|} \sum_{v \in \mathcal{V}_i} \nabla f(y_v, z_v^{(L)})$ is the minibatch gradient by the Exact estimator. Finally, $\nabla \mathcal{L}(W_i) = \frac{1}{|\mathcal{V}_{\mathcal{L}}|} \sum_{v \in \mathcal{V}_{\mathcal{L}}} \nabla f(y_v, z_{i,v}^{(L)})$ is the exact full-batch gradient.

The following lemma bounds the bias of the gradients by the CV estimator. Intuitively, there is a sequence of slow-changing model parameters (W_i) , where

W_i is the model at the i -th iteration. At each iteration i we use GCN with CV and GCN with Exact estimator to compute the activations for the minibatch \mathcal{V}_i , and update the corresponding history. After L epochs, we compute the gradient by backpropagating through CV's predictions on the minibatch of nodes \mathcal{V}_i . The gradient $g_{CV,i}(W_i)$ is a random variable of both the stochastic propagation matrix \hat{P}_i and the minibatch \mathcal{V}_i . But the expectation of the gradient w.r.t. \hat{P}_i and \mathcal{V}_i , $\mathbb{E}_{\hat{P}_i, \mathcal{V}_i} g_{CV,i}(W_i)$, is close to the full-batch gradient by the Exact estimator $\nabla \mathcal{L}(W_i)$, i.e., the gradient is close to be unbiased.

To study the gradient, we need the backpropagation rules of the networks. Let $f_v = f(y_v, z_v^{(L)})$ and $f_{CV,i,v} = f(y_z, z_{CV,i,v}^{(L)})$, we first derive the backpropagation rule for the exact algorithm. Differentiating both sides of Eq. (1), we have:

$$\begin{aligned} \nabla_{H^{(l)}} f_v &= P^\top \nabla_{Z^{(l+1)}} f_v W^{(l)\top} & l = 1, \dots, L-1 \\ \nabla_{Z^{(l)}} f_v &= \sigma'(Z^{(l)}) \circ \nabla_{H^{(l)}} f_v & l = 1, \dots, L-1 \\ \nabla_{W^{(l)}} f_v &= (PH^{(l)})^\top \nabla_{Z^{(l+1)}} f_v & l = 0, \dots, L-1. \end{aligned} \quad (5)$$

Similarly, differentiating both sides of Eq. (5), we have

$$\begin{aligned} \nabla_{H_{CV}^{(l)}} f_{CV,v} &= \hat{P}^{(l)} \nabla_{Z_{CV}^{(l+1)}} f_{CV,v} W^{(l)\top} & l = 1, \dots, L-1 \\ \nabla_{Z_{CV}^{(l)}} f_{CV,v} &= \sigma'(Z_{CV}^{(l)}) \circ \nabla_{H_{CV}^{(l)}} f_{CV,v} & l = 1, \dots, L-1 \\ \nabla_{W^{(l)}} f_{CV,v} &= (\hat{P}^{(l)} H_{CV}^{(l)})^\top \nabla_{Z_{CV}^{(l+1)}} f_{CV,v} & l = 0, \dots, L-1. \end{aligned} \quad (6)$$

Lemma 2. Assume $\sigma(\cdot)$ and $\nabla_z f(y, z)$ are ρ -Lipschitz, $\|\nabla_z f(y, z)\|_\infty \leq B$. Given a fixed dataset X and a sequence of T weights and stochastic propagation matrices $(W_i, \hat{P}_i)_{i=1}^T$, s.t.,

1. $\|W_i\|_\infty \leq B$, $\|\hat{P}_i\|_\infty \leq B$, and $\|\sigma'(Z_{CV,i})\|_\infty \leq B$,
2. $\|W_i - W_j\|_\infty < \epsilon, \forall i, j$,

let $P = \mathbb{E} \hat{P}_i$. If at time i we feed (X, W_i, \hat{P}_i) to a GCN with CV estimator to evaluate the prediction for nodes in the minibatch \mathcal{V}_i ,

$$Z_{CV,i}^{(l+1)} = \left(\hat{P}_i^{(l)} (H_{CV,i}^{(l)} - \bar{H}_{CV,i}^{(l)}) + P \bar{H}_{CV,i}^{(l)} \right) W_i^{(l)}, \quad H_{CV,i}^{(l+1)} = \sigma(Z_{CV,i}^{(l+1)}).$$

where $\bar{H}_{CV,i}^{(l)}$ is the maintained history at time i , and (X, W_i, P) to a GCN with exact estimator

$$Z_i^{(l+1)} = PH_i^{(l)} W_i^{(l)}, \quad H_i^{(l+1)} = \sigma(Z_i^{(l+1)}),$$

then there exists K that depends on C, B and ρ s.t.,

$$\left\| \mathbb{E}_{\hat{P}_i, \mathcal{V}_i} (W_i) - \nabla \mathcal{L}(W_i) \right\|_\infty \leq K\epsilon, \forall i > LI.$$

Proof. By Lipschitz continuity of $\nabla_z f(y, z)$ and Lemma 1, there exists \dot{K} , for all $\hat{P}_i = (\hat{P}_i^{(0)}, \dots, \hat{P}_i^{(L-1)})$

$$\begin{aligned} & \left\| \nabla_{z_{CV,v}^{(L)}} f_{CV,v} - \nabla_{z_v^{(L)}} f_v \right\|_\infty \\ & \leq \rho \left\| z_{CV,v}^{(L)} - z_v^{(L)} \right\|_\infty \\ & \leq \rho \dot{K} \epsilon, \\ & \left\| \sigma'(Z_{CV}^{(l)}) - \sigma'(Z^{(l)}) \right\|_\infty \leq \rho \dot{K} \epsilon \end{aligned} \quad (7)$$

We prove by induction that there exists K_l , s.t., $\forall l \in [L]$,

$$\left\| \mathbb{E}_{\hat{P}(\geq l)} \nabla_{z_{CV,v}^{(l)}} f_{CV,v} - \nabla_{z_v^{(l)}} f_v \right\|_\infty \leq K_l \epsilon, \quad \forall \hat{P}^{(0)}, \dots, \hat{P}^{(l-1)}, \quad (8)$$

where $\hat{P}(\geq l) = (\hat{P}^{(l)}, \dots, \hat{P}^{(L-1)})$. By Eq. (7) the statement holds for $l = L$, where $K_L = \rho \dot{K}$. If the statement holds for $l + 1$, i.e.,

$$\left\| \mathbb{E}_{\hat{P}(\geq l+1)} \nabla_{z_{CV,v}^{(l+1)}} f_{CV,v} - \nabla_{z_v^{(l+1)}} f_v \right\|_\infty \leq K_l \epsilon, \quad \forall \hat{P}^{(0)}, \dots, \hat{P}^{(l)},$$

then by Eq. (5, 6),

$$\begin{aligned} & \left\| \mathbb{E}_{\hat{P}(\geq l)} \nabla_{z_{CV,v}^{(l)}} f_{CV,v} - \nabla_{z_v^{(l)}} f_v \right\|_\infty \\ & = \left\| \mathbb{E}_{\hat{P}(\geq l)} \sigma'(Z_{CV}^{(l)}) \circ \hat{P}^{(l)} \nabla_{Z_{CV}^{(l+1)}} f_{CV,v} - \sigma'(Z^{(l)}) \circ P^\top \nabla_{Z^{(l+1)}} f_v \right\|_\infty \\ & = \left\| \mathbb{E}_{\hat{P}(\geq l)} \left\{ \left[\sigma'(Z_{CV}^{(l)}) \circ \hat{P}^{(l)} \nabla_{Z_{CV}^{(l+1)}} f_{CV,v} \right] - \sigma'(Z^{(l)}) \circ P^\top \nabla_{Z^{(l+1)}} f_v \right\} \right\|_\infty \\ & \leq \left\| \mathbb{E}_{\hat{P}(\geq l)} \left\{ \left[\left(\sigma'(Z_{CV}^{(l)}) - \sigma'(Z^{(l)}) \right) \circ \hat{P}^{(l)} \nabla_{Z_{CV}^{(l+1)}} f_{CV,v} \right] \right\} \right\|_\infty \\ & \quad + \mathbb{E}_{\hat{P}^{(l)}} \left\| \mathbb{E}_{\hat{P}(\geq l+1)} \left[\sigma'(Z^{(l)}) \circ \hat{P}^{(l)} \left(\nabla_{Z_{CV}^{(l+1)}} f_{CV,v} - \nabla_{Z^{(l+1)}} f_v \right) \right] \right\|_\infty \\ & \quad + \left\| \mathbb{E}_{\hat{P}(\geq l)} \left\{ \left[\sigma'(Z^{(l)}) \circ \left(\hat{P}^{(l)} - P^\top \right) \nabla_{Z^{(l+1)}} f_v \right] \right\} \right\|_\infty \\ & \leq \mathbb{E}_{\hat{P}(\geq l)} \left[\left\| \sigma'(Z_{CV}^{(l)}) - \sigma'(Z^{(l)}) \right\|_\infty \left\| \hat{P}^{(l)} \nabla_{Z_{CV}^{(l+1)}} f_{CV,v} \right\|_\infty \right] \\ & \quad + \mathbb{E}_{\hat{P}^{(l)}} \left\| \sigma'(Z^{(l)}) \right\|_\infty \mathbb{E}_{\hat{P}(\geq l+1)} \left\| \hat{P}^{(l)} \right\|_\infty \mathbb{E}_{\hat{P}(\geq l+1)} \left\| \nabla_{Z_{CV}^{(l+1)}} f_{CV,v} - \nabla_{Z^{(l+1)}} f_v \right\|_\infty \\ & \quad + 0 \\ & \leq \rho \dot{K} \epsilon B^2 C^2 + B^2 C^2 K_{l+1} \epsilon \\ & = K_l \epsilon, \end{aligned}$$

where $K_l = B^2 C^2 (\rho \dot{K} + K_{l+1})$. By induction, Eq. (8) holds. Similarly, we can show that there exists K , s.t.,

$$\left\| \mathbb{E}_{\hat{P}} \nabla_{W^{(l)}} f_{CV,v} - \nabla_{W^{(l)}} f_v \right\|_\infty < K \epsilon, \quad \forall l \in [L-1].$$

Therefore,

$$\begin{aligned}
& \left\| \mathbb{E}_{\mathcal{V}_i, \hat{P}_i} g_{CV,i}(W_i) - \nabla \mathcal{L}(W_i) \right\|_\infty \\
&= \left\| \mathbb{E}_{v \in \mathcal{V}, \hat{P}_i} g_{CV,v}(W_i) - \mathbb{E}_{v \in \mathcal{V}} g_v(W_i) \right\|_\infty \\
&\leq \mathbb{E}_{\mathcal{V}_i} \left\| \mathbb{E}_{\hat{P}_i} g_{CV,v}(W_i) - g_v(W_i) \right\|_\infty \\
&\leq \mathbb{E}_{\mathcal{V}_i} \max_t \left\| \mathbb{E}_{\hat{P}} \nabla_{W^{(t)}} f_{CV,v} - \nabla_{W^{(t)}} f_v \right\|_\infty \\
&\leq K\epsilon.
\end{aligned}$$

□

C.4 Proof of Theorem 2

Theorem 2. Assume that (1) the activation $\sigma(\cdot)$ is ρ -Lipschitz, (2) the gradient of the cost function $\nabla_z f(y, z)$ is ρ -Lipschitz and bounded, (3) $\|g_{CV,\mathcal{V}}(W)\|_\infty$, $\|g(W)\|_\infty$, and $\|\nabla \mathcal{L}(W)\|_\infty$ are all bounded by $G > 0$ for all \hat{P}, \mathcal{V} and W . (4) The loss $\mathcal{L}(W)$ is ρ -smooth, i.e., $|\mathcal{L}(W_2) - \mathcal{L}(W_1) - \langle \nabla \mathcal{L}(W_1), W_2 - W_1 \rangle| \leq \frac{\rho}{2} \|W_2 - W_1\|_F^2 \forall W_1, W_2$, where $\langle A, B \rangle = \text{tr}(A^\top B)$ is the inner product of matrix A and matrix B . (5) The loss $\mathcal{L}(W) \geq \mathcal{L}_*$ is bounded below. Then, there exists $K > 0$, s.t., $\forall N > LI$, if we run SGD for $R \leq N$ iterations, where R is chosen uniformly from $[N]_+$, we have

$$\mathbb{E}_R \|\nabla \mathcal{L}(W_R)\|_F^2 \leq 2 \frac{\mathcal{L}(W_1) - \mathcal{L}_* + K + \rho K}{\sqrt{N}},$$

for the updates $W_{i+1} = W_i - \gamma g_{CV,i}(W_i)$ and the step size $\gamma = \min\{\frac{1}{\rho}, \frac{1}{\sqrt{N}}\}$.

Proof. This proof is a modification of [2], but using biased stochastic gradients instead. We assume the algorithm is already warmed-up for LI steps with the initial weights W_0 , so that Lemma 2 holds for step $i > 0$. Denote $\delta_i = g_{CV,i}(W_i) - \nabla \mathcal{L}(W_i)$. By smoothness we have

$$\begin{aligned}
\mathcal{L}(W_{i+1}) &\leq \mathcal{L}(W_i) + \langle \nabla \mathcal{L}(W_i), W_{i+1} - W_i \rangle + \frac{\rho}{2} \gamma^2 \|g_{CV,i}(W_i)\|_F^2 \\
&= \mathcal{L}(W_i) - \gamma \langle \nabla \mathcal{L}(W_i), g_{CV,i}(W_i) \rangle + \frac{\rho}{2} \gamma^2 \|g_{CV,i}(W_i)\|_F^2 \\
&= \mathcal{L}(W_i) - \gamma \langle \nabla \mathcal{L}(W_i), \delta_i \rangle - \gamma \|\nabla \mathcal{L}(W_i)\|^2 + \frac{\rho}{2} \gamma^2 \left[\|\delta_i\|^2 + \|\nabla \mathcal{L}(W_i)\|_F^2 + 2\langle \delta_i, \nabla \mathcal{L}(W_i) \rangle \right] \\
&= \mathcal{L}(W_i) - (\gamma - \rho\gamma^2) \langle \nabla \mathcal{L}(W_i), \delta_i \rangle - (\gamma - \frac{\rho\gamma^2}{2}) \|\nabla \mathcal{L}(W_i)\|_F^2 + \frac{\rho}{2} \gamma^2 \|\delta_i\|_F^2.
\end{aligned} \tag{9}$$

For each i , consider the sequence of $LI + 1$ weights W_{i-LI}, \dots, W_i .

$$\begin{aligned} \max_{i-LI \leq j, k \leq i} \|W_j - W_k\|_\infty &\leq \sum_{j=i-LI}^{i-1} \|W_j - W_{j+1}\|_\infty \\ &= \sum_{j=i-LI}^{i-1} \gamma \|g_{CV}(W_j)\|_\infty \leq \sum_{j=i-LI}^{i-1} \gamma G = LIG\gamma. \end{aligned}$$

By Lemma 2, there exists $\dot{K} > 0$, s.t.

$$\left\| \mathbb{E}_{\hat{P}, \nu_B} \delta_i \right\|_\infty = \left\| \mathbb{E}_{\hat{P}, \nu_B} g_{CV}(W_i) - \nabla \mathcal{L}(W_i) \right\|_\infty \leq \dot{K} LIG\gamma, \quad \forall i > 0.$$

Assume that W is D -dimensional,

$$\begin{aligned} \mathbb{E}_{\hat{P}, \nu_B} \langle \nabla \mathcal{L}(W_i), \delta_i \rangle &\leq D \|\nabla \mathcal{L}(W_i)\|_\infty \left\| \mathbb{E}_{\hat{P}, \nu_B} \delta_i \right\|_\infty \leq \dot{K} LIDG^2\gamma \leq K\gamma, \\ \mathbb{E}_{\hat{P}, \nu_B} \|\delta_i\|_F^2 &\leq D \|g_{CV,i}(W_i)\|_\infty + D \|\nabla \mathcal{L}(W_i)\|_\infty \leq 2DG^2 \leq K, \end{aligned}$$

where $K = \max\{\dot{K}LIDG^2, 2DG^2\}$. Taking $\mathbb{E}_{\hat{P}, \nu_B}$ to both sides of Eq. 9 we have

$$\mathcal{L}(W_{i+1}) \leq \mathcal{L}(W_i) + (\gamma - \rho\gamma^2)K\gamma - (\gamma - \frac{\rho\gamma^2}{2}) \|\nabla \mathcal{L}(W_i)\|_F^2 + \rho K\gamma^2/2.$$

Summing up the above inequalities and re-arranging the terms, we obtain,

$$\begin{aligned} &(\gamma - \frac{\rho\gamma^2}{2}) \sum_i \|\nabla \mathcal{L}(W_i)\|_F^2 \\ &\leq \mathcal{L}(W_1) - \mathcal{L}_* + KN(\gamma - \rho\gamma^2)\gamma + \frac{\rho K}{2} N\gamma^2. \end{aligned}$$

Dividing both sides by $N(\gamma - \frac{\rho\gamma^2}{2})$, and take $\gamma = \min\{\frac{1}{\rho}, \frac{1}{\sqrt{N}}\}$

$$\begin{aligned} &\mathbb{E}_{R \sim P_R} \|\nabla \mathcal{L}(W_R)\|_F^2 \\ &\leq 2 \frac{\mathcal{L}(W_1) - \mathcal{L}_* + KN(\gamma - \rho\gamma^2)\gamma + \frac{\rho K}{2} N\gamma^2}{N\gamma(2 - \rho\gamma)} \\ &\leq 2 \frac{\mathcal{L}(W_1) - \mathcal{L}_* + KN(\gamma - \rho\gamma^2)\gamma + \frac{\rho K}{2} N\gamma^2}{N\gamma} \\ &\leq 2 \frac{\mathcal{L}(W_1) - \mathcal{L}_*}{N\gamma} + K\gamma(1 - \rho\gamma) + \rho K\gamma \\ &\leq 2 \frac{\mathcal{L}(W_1) - \mathcal{L}_*}{\sqrt{N}} + K\gamma + \rho K/\sqrt{N} \\ &\leq 2 \frac{\mathcal{L}(W_1) - \mathcal{L}_* + K + \rho K}{\sqrt{N}}. \end{aligned}$$

Particularly, when $N \rightarrow \infty$, we have $\mathbb{E}_{R \sim P_R} \|\nabla \mathcal{L}(W_R)\|_F^2 = 0$, which implies that the gradient is asymptotically unbiased. \square

C.5 Generalizing to Graph Attention Networks

Our Theorem 2 can generalize to graph attention networks (GAT) [6]. GAT updates can be written as

$$Z^{(l+1)} = P(H^{(l)}, W^{(l)})H^{(l)}W^{(l)}, \quad H^{(l+1)} = \sigma(Z^{l+1}),$$

the difference between GAT and GCN (Eq. 1) is that the propagation matrix P is now a function of activations and network weights instead of a constant. We define the variance reduced stochastic update as

$$Z^{(l+1)} = \left(\hat{P}^{(l)}(\bar{H}^{(l)}, W^{(l)})(H^{(l)} - \bar{H}^{(l)}) + P(\bar{H}^{(l)}, W^{(l)})\bar{H}^{(l)} \right) W^{(l)},$$

where we approximate $H^{(l)}$ with $\bar{H}^{(l)}$ when computing the propagation matrix. We can still bound the gradient by extending Lemma 2, and prove the convergence.

For GraphSAGE-pool and GraphSAGE-LSTM, our algorithm does not directly apply. Take GraphSAGE-pool as an example, it defines

$$z_{ud}^{(l+1)} = \max_{v \in \mathbf{n}(u)} \left(h_v^{(l)} W_{:,d}^{(l)} \right),$$

whose gradient is

$$\frac{\partial z_{ud}^{(l+1)}}{\partial h_{v,d'}^{(l)}} = \begin{cases} w_{d'd}^{(l)} & \text{if } v \text{ maximizes } h_v^{(l)} W_{:,d}^{(l)}, \\ 0 & \text{otherwise.} \end{cases}$$

It is unclear how to obtain an unbiased stochastic approximation of this gradient. We leave this as an open problem to study.

D Pseudocode

As mentioned in Sec. 3.3, an iteration of our algorithm consists the following operations:

1. Randomly select a minibatch $\mathcal{V}_B \in \mathcal{V}_L$ of nodes;
2. Build a computation graph that only contains the activations $h_v^{(l)}$ and $\bar{h}_v^{(l)}$ needed for the current minibatch;
3. Get the predictions by forward propagation as Eq. (6) in the main text;
4. Get the gradients by backward propagation, and update the parameters by SGD;
5. Update the historical activations.

For step 2, we construct the receptive fields $\mathbf{r}^{(l)}$ and stochastic propagation matrices $\hat{P}^{(l)}$ as Alg. 1.

Algorithm 1 Constructing the receptive fields and random propagation matrices.

```

 $\mathbf{r}^{(L)} \leftarrow \mathcal{V}_B$ 
for layer  $l \leftarrow L - 1$  to 0 do
   $\mathbf{r}^{(l)} \leftarrow \emptyset$ 
   $\hat{P}^{(l)} \leftarrow \mathbf{0}$ 
  for each node  $u \in \mathbf{r}^{(l+1)}$  do
     $\mathbf{r}^{(l)} \leftarrow \mathbf{r}^{(l)} \cup \{u\}$ 
     $\hat{P}_{uu}^{(l)} \leftarrow \hat{P}_{uu}^{(l)} + P_{uu}n(u)/D^{(l)}$ 
    for  $D^{(l)} - 1$  random neighbors  $v \in \mathbf{n}(u)$  do
       $\mathbf{r}^{(l)} \leftarrow \mathbf{r}^{(l)} \cup \{v\}$ 
       $\hat{P}_{uv}^{(l)} \leftarrow \hat{P}_{uv}^{(l)} + P_{uv}n(u)/D^{(l)}$ 
    end for
  end for
end for

```

D.1 Training with the CV estimator

Alg. 2 depicts the training algorithm using the CV estimator. We perform forward propagation according to Eq. (6), compute the stochastic gradient, and then update the historical activations $\bar{H}^{(l)}$ for all the nodes in $\mathbf{r}^{(l)}$. Let $W = (W^{(0)}, \dots, W^{(L-1)})$ be all the trainable parameters, the gradient $\nabla_W \mathcal{L}$ is computed automatically by frameworks such as TensorFlow.

D.2 Training with the CVD estimator

Training with the CVD estimator is similar with the CV estimator, except it runs two versions of the network, with and without dropout, to compute the samples H and their mean μ of the activation. The matrix $\bar{P}_{uv}^{(l)} = \hat{P}_{uv}^{(l)} / \sqrt{n(v)}$, where $n(v)$ is the degree of node v .

E Experiment setup

In this sections we describe the details of our model architectures. We use the Adam optimizer [4] with learning rate 0.01.

- Citeseer, Cora, PubMed and NELL: We use the same architecture as [5]: two graph convolution layers with one linear layer per graph convolution layer. We use 32 hidden units, 50% dropout rate and 5×10^{-4} L2 weight decay for Citeseer, Cora and PubMed and 64 hidden units, 10% dropout rate and 10^{-5} L2 weight decay for NELL.
- PPI and Reddit: We use the mean pooling architecture GraphSAGE-mean proposed by [3]. We use two linear layers per graph convolution layer. We set weight decay as zero, dropout rate as 20%, and adopt layer

Algorithm 2 Training with the CV algorithm

for each minibatch $\mathcal{V}_B \subset \mathcal{V}$ **do**
 Compute the receptive fields $\mathbf{r}^{(l)}$ and stochastic propagation matrices $\hat{P}^{(l)}$
 as Alg. 1.
 (Forward propagation)
 for each layer $l \leftarrow 0$ to $L - 1$ **do**
 $Z^{(l+1)} \leftarrow \left(\hat{P}^{(l)}(H^{(l)} - \bar{H}^{(l)} + P\bar{H}^{(l)}) \right) W^{(l)}$
 $H^{(l+1)} \leftarrow \sigma(Z^{(l+1)})$
 end for
 Compute the loss $\mathcal{L} = \frac{1}{|\mathcal{V}_B|} \sum_{v \in \mathcal{V}_B} f(y_v, Z_v^{(L)})$
 (Backward propagation)
 $W \leftarrow W - \gamma_i \nabla_W \mathcal{L}$
 (Update historical activations)
 for each layer $l \leftarrow 0$ to $L - 1$ **do**
 for each node $v \in \mathbf{r}^{(l)}$ **do**
 $\bar{h}_v^{(l)} \leftarrow h_v^{(l)}$
 end for
 end for
end for

normalization [1] after each linear layer. We use 512 hidden units for PPI and 128 hidden units for Reddit. We find that our architecture can reach 97.8% testing micro-F1 on the PPI dataset, which is significantly higher than 59.8% reported by [3]. We find the improvement is from wider hidden layer, dropout and layer normalization.

F Experiment for 3-layer GCNs

We test 3-layer GCNs on the Reddit dataset. The settings are the same with 2-layer GCNs in Sec. 6.2. To ensure M1+PP can run in a reasonable amount of time, we subsample the graph so that the maximum degree is 10. The convergence result is shown as Fig. 1, where the conclusion is similar with the two-layer models: CVD+PP is the best-performing approximate algorithm, followed by CV+PP, and then NS+PP and NS. The time consumption to reach 0.94 testing accuracy is shown in Table 1.

G Correlation between node activations

In our analysis of the variance for the CVD estimator in Sec. 5.2, we assume that the activations for different nodes are uncorrelated, i.e., $\text{Cov}_M [h_u^{(l)}, h_v^{(l)}] = 0$, for all $u \neq v$, where M is the dropout mask. We show the rationale behind this assumption in this section. For 2-layer GCNs, the activations are indeed

Algorithm 3 Training with the CVD algorithm

for each minibatch $\mathcal{V}_B \subset \mathcal{V}$ **do**
 Compute the receptive fields $\mathbf{r}^{(l)}$ and stochastic propagation matrices $\hat{P}^{(l)}$
 as Alg. 1.
 (Forward propagation)
 for each layer $l \leftarrow 0$ to $L - 1$ **do**
 $U \leftarrow \left(\bar{P}^{(l)}(H^{(l)} - \mu^{(l)}) + \hat{P}^{(l)}(\mu^{(l)} - \bar{\mu}^{(l)}) + P\bar{H}^{(l)} \right)$
 $H^{(l+1)} \leftarrow \sigma(\text{Dropout}_p(U)W^{(l)})$
 $\mu^{(l+1)} \leftarrow \sigma(UW^{(l)})$
 end for
 Compute the loss $\mathcal{L} = \frac{1}{|\mathcal{V}_B|} \sum_{v \in \mathcal{V}_B} f(y_v, H_v^{(L)})$
 (Backward propagation)
 $W \leftarrow W - \gamma_i \nabla_W \mathcal{L}$
 (Update historical activations)
 for each layer $l \leftarrow 0$ to $L - 1$ **do**
 for each node $v \in \mathbf{r}^{(l)}$ **do**
 $\bar{h}_v^{(l)} \leftarrow h_v^{(l)}$
 end for
 end for
end for

Table 1: Time to reach 0.95 testing accuracy.

Alg.	Valid. acc.	Epochs	Time (s)	Sparse GFLOP	Dense TFLOP
Exact	0.940	3.0	199	306	11.7
NS	0.940	24.0	148	33.6	9.79
NS+PP	0.940	12.0	68	2.53	4.89
CV+PP	0.940	5.0	32	8.06	2.04
CVD+PP	0.940	5.0	36	16.1	4.08

independent, and the correlation is still weak for deeper GCNs due to the sparsity of our sampled graph.

G.1 Results for 2-layer GCNs

For a 2-layer GCN with the first layer pre-processed, the activations of nodes are independent. Suppose we want to compute the prediction for a node on the second layer. Without loss of generality, assume that we want to compute $z_1^{(2)}$, and the neighbors of node 1 are $1, \dots, D$. The activation $h_v^{(1)} = \sigma\left((M_v \circ u_v^{(0)})W^{(0)}\right)$, where $u_v^{(0)} = (PH^{(0)})_v$ is a random variable with respect to M_v , $M_v \sim \text{Bernoulli}(p)$ is the dropout mask. We show that $h_v^{(1)}$ and $h_{v'}^{(1)}$ are independent, for $v \neq v'$ by the following lemma.

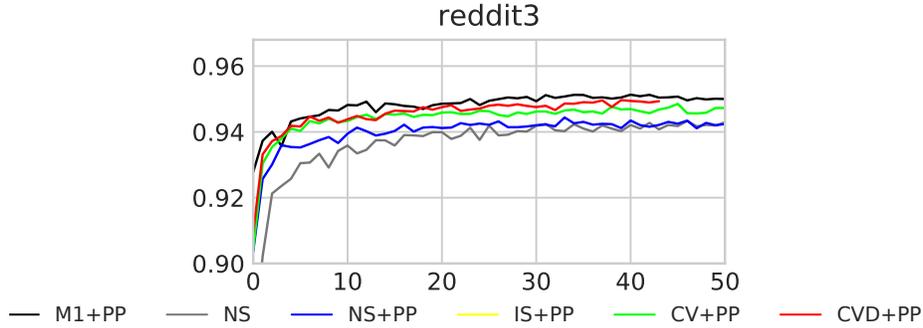


Figure 1: Comparison of validation accuracy with respect to number of epochs for 3-layer GCNs.

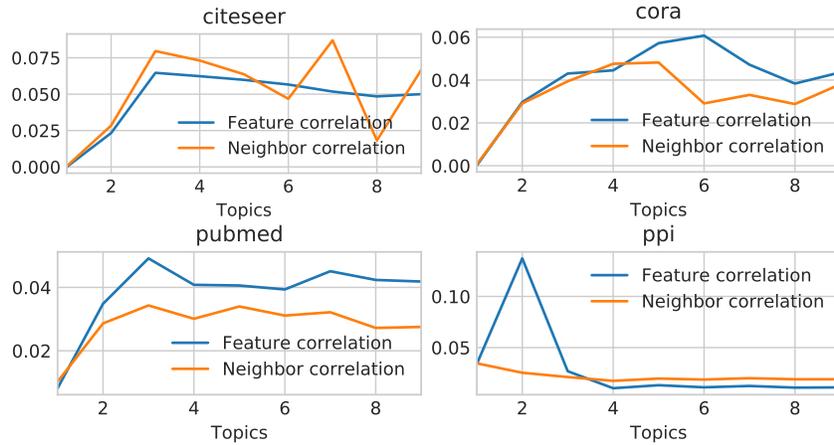


Figure 2: Average feature and neighbor correlations in a 10-layer GCN.

Lemma 3. *If a and b are independent random variables, then their transformations $f_1(a)$ and $f_2(b)$ are independent.*

Because for any event A and B , $P(f_1(a) \in f_1(A), f_2(b) \in f_2(B)) = P(a \in A, b \in B) = P(a \in A)P(b \in B) = P(f_1(a) \in f_1(A))P(f_2(b) \in f_2(B))$, where $f_1(A) = \{f_1(a)|a \in A\}$ and $f_2(B) = \{f_2(b)|b \in B\}$.

Let $h_v^{(1)} = f_1(M_v) := \sigma\left((M_v \circ u_v^{(0)})W^{(0)}\right)$ and $h_{v'}^{(1)} = f_1(M_{v'}) := \sigma\left((M_{v'} \circ u_{v'}^{(0)})W^{(0)}\right)$,

because M_v and $M_{v'}$ are independent Bernoulli random variables, $h_v^{(1)}$ and $h_{v'}^{(1)}$ are independent.

The result can be further generalized to deeper models. If the receptive fields of two nodes does not overlap, they should be independent.

G.2 Empirical results for deeper GCNs

Because we only sample two neighbors per node, the sampled subgraph is very close to a graph with all its nodes isolated, which reduces to the MLP case that [7] discuss.

We empirically study the correlation between feature dimensions and neighbors. The definition of the correlation between feature dimensions is the same with [7]. For each node v on layer l , we compute the correlation between each feature dimension of $h_v^{(l)}$

$$\begin{aligned} \text{Cov}_{ij}^{(l,v)} &:= \mathbb{C}[h_{vi}^{(l)}, h_{vj}^{(l)}] \\ \text{Corr}_{ij}^{(l,v)} &:= \frac{\text{Cov}_{ij}^{(l,v)}}{\sqrt{\text{Cov}_{ii}^{(l,v)}} \sqrt{\text{Cov}_{jj}^{(l,v)}}}, \end{aligned}$$

where i and j are the indices for different hidden dimensions, and $\mathbb{C}[X, Y] = \mathbb{E}[(X - \mathbb{E}X)(Y - \mathbb{E}Y)]$ is the covariance between two random variables X and Y . We approximate $\text{Cov}_{ij}^{(l,v)}$ with 1,000 samples of the activations $h_{vi}^{(l)}$ and $h_{vj}^{(l)}$, by running the forward propagation 1,000 times with different dropout masks. We define the *average feature correlation on layer l* to be $\text{Cov}_{ij}^{(l,v)}$ averaged by the nodes v and dimension pairs $i \neq j$.

To compute the correlation between neighbors, we treat each feature dimension separately. For each layer $l + 1$, node v , and dimension d , we compute the correlation matrix of all the activations $\{h_{id}^{(l)} | i \in \bar{\mathbf{n}}^{(l)}(v)\}$ that are needed by $h_{vd}^{(l+1)}$, where $\bar{\mathbf{n}}^{(l)}(v) = \{i | \hat{P}_{vi}^{(l)} \neq 0\}$ is the set of subsampled neighbors for node v :

$$\begin{aligned} \text{Cov}_{ij}^{(l,v,d)} &:= \mathbb{C}[h_{id}^{(l)}, h_{jd}^{(l)}] \\ \text{Corr}_{ij}^{(l,v,d)} &:= \frac{\text{Cov}_{ij}^{(l,v,d)}}{\sqrt{\text{Cov}_{ii}^{(l,v,d)}} \sqrt{\text{Cov}_{jj}^{(l,v,d)}}}, \end{aligned}$$

where the indices $i, j \in \bar{\mathbf{n}}^{(l)}(v)$. Then, we compute the average correlation of all pairs of neighbors $i \neq j$.

$$\text{AvgCorr}^{(l,v,d)} := \frac{1}{|\bar{\mathbf{n}}^{(l)}(v)| (|\bar{\mathbf{n}}^{(l)}(v)| - 1)} \sum_{i \neq j} \text{Corr}_{ij}^{(l,v,d)},$$

and define the *average neighbor correlation on layer l* as $\text{AvgCorr}^{(l,v,d)}$ averaged over all the nodes v and dimensions d .

We report the average feature correlation and the average neighbor correlation per layer, on the Citeseer, Cora, PubMed and PPI datasets. These quantities are too expensive to compute for NELL and Reddit. On each dataset, we train a GCN with 10 graph convolution layers until early stopping criteria is met, and compute the average feature correlation and the average neighbor correlation for

Dataset	Citeseer	Cora	PubMed	NELL	Reddit	PPI
w.o. dropout	70.2 ± .6	78.3 ± .4	77.7 ± .2	64.5 ± .7	95.6 ± .07	90.6 ± .6
w. dropout	69.9 ± .6	78.7 ± .3	78.3 ± .8	65.0 ± .1	96.5 ± .05	97.3 ± .03

Table 2: Validating accuracy / micro-F1 for models with or without dropout.

layer 1 to 9. We are not interested in the correlation on layer 10 because there are no more graph convolutional layers after it. The result is shown as Fig. 2. As analyzed in Sec. G.1, the average neighbor correlation is close to zero on the first layer, but it is not exactly zero due to the finite sample size for computing the empirical covariance. There is no strong tendency of increased correlation as the number of layers increases, after the third layer. The average neighbor correlation and the average feature correlation remain on the same order of magnitude, so bringing correlated neighbors does not make the activations much more correlated than the MLP case [7]. Finally, both correlations are much smaller than one.

H Effect of Dropout

We compare the models with or without dropout as Table 2.

References

- [1] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- [2] Saeed Ghadimi and Guanghui Lan. Stochastic first-and zeroth-order methods for nonconvex stochastic programming. *SIAM Journal on Optimization*, 23(4):2341–2368, 2013.
- [3] Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. In *Advances in Neural Information Processing Systems*, pages 1025–1035, 2017.
- [4] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2014.
- [5] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *ICLR*, 2017.
- [6] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks. In *ICLR*, 2018.
- [7] Sida Wang and Christopher Manning. Fast dropout training. In *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*, pages 118–126, 2013.