ON BATCH SIZE SELECTION FOR STOCHASTIC TRAINING FOR GRAPH NEURAL NETWORKS

Anonymous authors
Paper under double-blind review

ABSTRACT

In recent years deep learning has become an important framework for supervised learning. It has been observed that stochastic gradient decent (SGD) method in deep learning networks performs well when the minibatch size is small. In this work, we focus on the importance of batch size selection in Graph Neural Networks (GNN). We provide theoretical analysis based on an estimator that considers the randomness arising from two consecutive layers in GNN, and suggest a guideline for picking the appropriate scale of the batch size. We complement our theoretical results with empirical experiments. We consider the following baseline methods: ClusterGCN, FastGCN, GraphSaint on the following datasets: Ogbn-products, Ogbn-arxiv, Reddit and Pubmed. We demonstrate that in contrast to conventional deep learning models, GNNs benefit from large batch sizes.

1 INTRODUCTION

Training large neural networks is often time consuming. In many real world scenarios training might take hours or even days to converge Radford et al. (2018); Devlin et al. (2018). As a consequence, the identification of strategies to reduce the training time while retaining accuracy is an important research objective. The most popular training algorithms for deep learning are Stochastic Gradient Descent (SGD) and its variants such as RMSProp or Adam Graves (2013); Kingma & Ba (2014). These algorithms work in an iterative manner, such that in each epoch, the data is first partitioned into *minibatches* and then weight updates are calculated using only the data in each mimibatch. It has been observed that the size of the minibatches plays a crucial role in the network's accuracy, generalization capability and converge time (Keskar et al. (2016); He et al. (2019); McCandlish et al. (2018); Radiuk (2017)).

The maximum size of the minibatch is limited by the available memory, but for typical deep learning tasks, practitioners have observed that small batch sizes, e.g., $\{4, 16, \ldots, 512\}$, lead to a better generalization performance and training efficiency. For Graph Neural Networks (GNNs) selecting the appropriate batch size remains more of a mystery, and to the best of our knowledge, there has been no published work that focuses on batch size selection for GNNs. The small batch size guidelines for conventional NNs do not carry over because the batches are used to approximate the graph aggregations or convolutions. The approximation error propagates and leads to a much more substantial variance in the gradients than is observed for NNs. In practice, based on released code, we see that implementations tend to either use the largest batch size that can fit into memory Li et al. (2020) or use a small batch size similar to those for non-graph settings Chen et al. (2018); Zou et al. (2019).

In this work, we explore the choice of batch size for graph neural networks. By means of a theoretical investigation, we develop guidelines for the choice of batch size that depend on the average degree and number of nodes of the graph. These guidelines lead to intermediate batch sizes, considerably larger than the small NN batch sizes but much smaller than the maximum size dictated by memory limits of a modern GPU. We provide empirical results that demonstrate that the batch sizes derived using our guidelines provide an excellent trade-off between training time and accuracy. Substantially smaller sizes may lead to faster convergence but reduce accuracy; using larger sizes can achieve similar accuracy but training take much longer to converge.

2 RELATED WORK

Graph Neural Networks (GNNs) have become increasingly popular in addressing graph-based tasks (Kipf & Welling, 2016; Hamilton et al., 2017; Defferrard et al., 2016; Gilmer et al., 2017; Ying et al., 2018). One major line of research aims to improve the expressiveness of GNNs via 1) advanced aggregation functions (Veličković et al., 2017; Monti et al., 2017; Liu et al., 2019; Qu et al., 2019; Pei et al., 2020) 2) deeper architecture (Li et al., 2019; 2020); and 3) adaptive graph structure (Li et al., 2018; Vashishth et al., 2019; Zhang et al., 2019). However, training a large-scale GNN model remains challenging because of the large memory consumption, long convergence time, and heavy computation (Chiang et al., 2019).

Full-batch gradient descent training scheme was commonly used in the earlier GNN research. While this is suitable for relatively small graphs, it requires storing all intermediate embeddings, which is not scalable for large graphs. The convergence can be slow since the parameters are updated only once per epoch. Hamilton et al. (2017) and Ying et al. (2018) proposed the training of GNNs with mini-batch stochastic gradient descent (SGD) methods. Mini-batch SGD training suffers from the neighborhood expansion and leads to time-complexity that grows exponential with respect to the GNN depth. To reduce the exponential complexity of receptive nodes, Chen et al. (2018), Huang et al. (2018), and Zou et al. (2019) proposed layer-wise sampling, where a fixed number of nodes are sampled in each layer. Importance sampling techniques were incorporated to reduce variance. Unfortunately the overhead of the iterative neighborhood sampling strategy is still significant and becomes worse as GNNs become progressively deeper.

Chiang et al. (2019) and Zeng et al. (2020) proposed graph-wise sampling to further improve the sampling efficiency. This can be viewed as a special case of layer-wise sampling where the same set of nodes is sampled across all layers. Chen et al. (2017) and Cong et al. (2020) proposed variance reduction stochastic training frameworks that maintain a cache for the intermediate embeddings of all nodes. This can improve convergence but results in large memory requirements, stretching the capabilities of GPUs when training over large graphs. Due to this drawback, we do not consider such approaches in this paper, but it is an intrguing direction for future work.

Most existing graph neural network papers do not clearly address how they set the batch size. Experimentally, we observe that batch size is a critical hyper-parameter and can significantly influence training time and test accuracy. The importance of the batch size has been recognized for non-graph deep learning models. Keskar et al. (2016), He et al. (2019) and Masters & Luschi (2018) have shown that smaller batch sizes, in the range $\{4, 16, \ldots, 512\}$, can achieve better generalization performance. The randomness of small batches proves beneficial. McCandlish et al. (2018) suggested that the batch size should be selected so that a balance is achieved between the "noise" and "signal" of the gradient. Radiuk (2017) showed that using larger batch sizes, of the order of 1024, can be beneficial when training convolutional neural network models. Gower et al. (2019), Alfarra et al. (2020), and Smith (2018) introduced adaptive batch size approaches to further improve the convergence rate and generalization performance. To the best of our knowledge, no existing work has directly addressed the selection of the batch size for stochastic training for graph neural networks, and the objective of this paper is to fill that gap and provide guidelines for the GNN setting.

3 PRELIMINARIES

We represent a graph $G = (V, E)$ with a set of nodes $V = \{v_1, \ldots, v_n\}$ and set of edges $E = \{e_1, \ldots, e_M\}$ by an adjacency matrix $A \in \mathbb{R}^{n \times n}$. For node $v \in V$, we let $N(v)$ be the set of neighbors of v. In addition, we associate each node v to a feature vector $x_v \in \mathbb{R}^{1 \times F}$, and let $X \in \mathbb{R}^{n \times F}$ be the corresponding feature matrix. Let D be the degree matrix of the graph G, where $D_{i,i} = \sum_j A_{i,j}$ and $D_{i,j} = 0$ if $i \neq j$. To ease the presentation, we use symbols such as $R \lesssim T$ to denote that there exists an absolute constant c such that $R \leq c \cdot T$.

3.1 GRAPH NEURAL NETWORK MODELS

Graph neural networks (GNNs) can be applied to node prediction, link prediction and graph prediction tasks. In this work, we focus on the *node prediction* task. We are given labels of nodes from a training set and we need to predict the labels for nodes in a testing set. One paradigm for

solving this problem is to learn representations for all the nodes and then map the representations to labels. Graph neural network aggregate the representations of neighbors into each node in order to integrate the graph structure into each node's representation. Specifically, let $H^l \in \mathbb{R}^{N \times F^l}$ be the representation for layer l, where the ith row h_i^l is the representation for node i at layer l and F^l is the dimension of the representation (H^0 is set as the original node features X). The forward propagation for hidden states are defined as:

$$H^{l+1} = \sigma(\tilde{H}^{l+1}) \qquad \text{and} \qquad \tilde{H}^{l+1} = \tilde{A} H^l W^l, \tag{1}$$

where $W^l \in \mathbb{R}^{F^l \times F^{l+1}}$ are trainable parameters, $\sigma(\cdot)$ is an activation function, and $\tilde{A}$ is a normalization of adjacency matrix A, e.g., the *random walk* normalization $\tilde{A} = D^{-1}A$, or the *symmetric* normalization $\tilde{A} = D^{-1/2}AD^{-1/2}$. The equations in (1) can be expressed for each node i as:

$$h_i^{l+1} = \sigma(\tilde{h}_i^{l+1}) \qquad \text{and} \qquad \tilde{h}_i^{l+1} = \sum_{j \in N(i)} \tilde{A}_{ij} h_j^l W^l. \tag{2}$$

3.2 Stochastic Training for GNN Models

Two roles of sampling in training GNNs: Label sampling and neighbor sampling. In conventional deep learning models, every sample in a mini-batch contributes independently to the approximated gradient. Including more samples thus reduces the variance of the gradient estimate by statistical power. However, in GNN models, we have two different concepts of sampling. First, we need to select a mini-batch of nodes to include in the training set and we call this *label sampling*. Since the representation of nodes also depends on neighbor nodes, the receptive field for each selected node grows exponentially as the number of layers increases. We need a mechanism to constrain the number of receptive neighbor nodes. We call this *neighbor sampling*. Existing frameworks use three main approaches to handle sampling for GNNs: *node-wise*, *layer-wise* and *graph-wise* sampling.

Node-wise sampling. Hamilton et al. (2017) and Ying et al. (2018) adopt a uniformly random sampling of the labels. For neighbor sampling, they recursively sample a certain number of neighbors for each layer. Specifically, to evaluate the aggregation for layer $l+1$, for each node i, S_i^l is sampled from the neighbors of node i, and equation (2) is evaluated as

$$\tilde{h}_i^{l+1} = \frac{|N(i)|}{|S_i^l|} \sum_{j \in S_i^l} A_{ij} h_j^l W^l. \tag{3}$$

$|S_i^l|$ is predefined to limit the sampled nodes, but the size of the receptive field for each included mini-batch node still grows exponentially as the number of layers increases.

Layer-wise sampling. Chen et al. (2018), Huang et al. (2018),and Zou et al. (2019) propose a different neighbor sampling method to reduce the number of receptive nodes. Nodes are sampled at the layer level instead of the node level. Importance sampling is adopted to reduce the variance of sampling and further improve convergence. Specifically, given the set of sampled nodes S^{l+1} in layer $l+1$, nodes in layer l are sampled with some probability distribution $q^l(i|S^{l+1})$, where i is the index of the node. For brevity, we denote the sampling distribution as $q^l(i)$. The forward propagation in equation (2) is defined as:

$$\tilde{h}_i^{l+1} = \frac{1}{|S^l|} \sum_{j \in S^l} \frac{A_{ij}}{q^l(i)} h_j^l W^l. \tag{4}$$

Note that by controlling $|S^l|$, the number of receptive nodes only grow linearly with respect to layer size. In this framework, uniformly random sampling is used for labels.

Graph-wise sampling. Zeng et al. (2020) and Chiang et al. (2019) propose graph-wise sampling. This can be regarded as a special case of layer-wise sampling, which uses the same set of nodes as both the sampled labels and sampled neighbors across all layers. Importance sampling is adopted in Zeng et al. (2020) and normalization is applied on both loss and neighbor aggregations to obtain an unbiased estimator for the gradients.

Since in practice, layer-wise sampling and graph-wise sampling are much more efficient in practice than node-wise sampling, we focus on the influence of the batch size for layer-wise sampling and

graph-wise sampling. For layer-wise sampling, the number of included label samples (the batch size) can differ from the number of neighbours sampled at each layer. In our analysis, we focus on the case where these values are equal; see the supplementary material for further discussion.

4 BATCH SIZE SELECTION: THEORETICAL ANALYSIS

When we use SGD training, we sample a small proportion of samples to approximate the true distribution of samples and estimate the gradients. In the non-graph setting, samples in the mini-batch contribute (approximately) independently to the estimates of gradients. In GNN models, the impact is more complicated, because node representations are derived using a sample of the neighborhood. If all nodes were included, a node's representation would be calculated based on its entire neighborhood; for small batch sizes, only a few neighbours are included, and the approximation of the aggregation can be very poor. This propagates through the layers leading to highly erroneous gradient estimates. In this section, we consider the selection of the batch size for layer-wise and graph-wise sampling in GNN training and derive guidelines. Our approach is to analyse a metric for an estimator which captures both the variance and compute time.

4.1 VARIANCE OF GRADIENT, VARIANCE ESTIMATOR

The essential goal of sampling in SGD is to approximate the gradient. We aim to obtain an unbiased approximation and minimize the variance of the gradients in each minibatch. Recall the definition of the activation of a node j from (2) and let L be a loss function. By the chain rule, the gradient with respect to the variables in layer l, when all nodes are included (i.e., no sampling), is:

$$\frac{\partial L}{\partial W^l} = \sum_{i \in V^{l+1}} \frac{\partial L}{\partial \tilde{h}_i^{l+1}} \frac{\partial \tilde{h}_i^{l+1}}{\partial W^l} = \sum_{i \in V^{l+1}} \frac{\partial L}{\partial \tilde{h}_i^{l+1}} \sum_{j \in N(i)} A_{ij} h_j^l, \tag{5}$$

where V^l denotes the receptive nodes in layer l. Under layer-wise sampling, the gradient can be expressed as:

$$\frac{\partial L}{\partial W^l} = \sum_{i \in S^{l+1}} \frac{\partial L}{\partial \tilde{h}_i^{l+1}} \sum_{j \in S^l} \frac{A_{ij}}{|S^l| q_j^l} h_j^l, \tag{6}$$

where S^l and S^{l+1} are the sets of sampled nodes for layer l and layer $l+1$, respectively. Analyzing the exact variance in general is difficult. Instead, we analyze intermediate estimators, with the view that these can act as a valuable proxy for the variance of gradients.

In most existing work, the neighbor aggregation terms $\sum_{j \in S^l} \frac{A_{ij}}{|S^l| q_j^l} h_j^l$ is used as the proxy estimator (Chen et al., 2018; Huang et al., 2018; Zou et al., 2019; Zeng et al., 2020). This proxy estimator does not adequately capture the correlation between layers. The variance of the gradients for the variables in layer l is highly related to the nodes sampled in both layer $l+1$ and layer l. To address this issue, we analyze a different estimator which considers the randomness arising from two consecutive layers. Specifically, we consider the analysis of the following estimator.

Definition 4.1 *Let $\mathrm{S}_1, \mathrm{S}_2 \subseteq V$ such that each vertex in V is selected to S_1 (respectively S_2) with probability p (respectively q). For weight matrix W, we define our estimator as:*

$$\xi = \frac{1}{|\mathrm{S}_1|} \sum_{v \in \mathrm{S}_1} \frac{\mathbf{1}_{N(v) \cap \mathrm{S}_2 \neq \emptyset}}{|\mathrm{S}_2 \cap N(v)|} \sum_{u \in N(v) \cap \mathrm{S}_2} \frac{\tilde{A}_{v,u} \cdot W \cdot x_u}{\alpha_v},$$

where $\alpha_v = 1 - (1-q)^{|N(v)|}$, and $\mathbf{1}_{\mathrm{Z}}$ is the indicator random variable for the event Z.

4.2 PSEUDO PRECISION RATE

In practice, one of the major purposes of sampling is to improve training efficiency and reduce the time for training the model. Existing methods only aim to reduce the variance of the estimators and do not take the computational cost into account Chen et al. (2018); Huang et al. (2018); Zou et al.

(2019); Zeng et al. (2020). When we evaluate the impact of batch size, a larger batch size will generate better approximation of gradient, but this comes at the cost of significantly more computation. Therefore, we need a better metric that can better capture the trade-off between variance reduction and computation cost. McCandlish et al. (2018) propose a metric that balances the noise scale and gradient value in each minibatch to determine the optimal batch size, but it does not explicitly model the computational cost in the metric. In the context of variance reduction for Monte Carlo sampling, Owen (2013) introduces an efficiency metric for an estimator. If there is a reference estimator that has compute time c_0 and achieves variance σ_0^2, then the *efficiency* of an alternative estimator with variance σ_1^2 and compute time c_1 is defined as $\frac{c_0\sigma_0^2}{c_1\sigma_1^2}$. Normalizing compute time so that the reference estimator satisfies $c_0\sigma_0^2 = 1$, we derive the metric $\frac{1}{c\sigma^2}$, which we call the *pseudo precision rate*:

Definition 4.2 *Let ξ be an estimator with computation cost $c > 0$ and variance $\sigma^2 > 0$, then the* pseudo precision rate *of ξ is defined as*

$$\rho(\xi) = \frac{1}{c\sigma^2}. \tag{7}$$

Intuitively, this metric characterizes how much we can reduce the variance per unit computation time. By maximizing the pseudo precision rate, we can achieve a balance between variance reduction and computational cost.

4.3 Guideline for Batch Size

We derive the guideline for selecting batch size by analyzing how the batch size influences the pseudo precision rate of the estimator ξ in Definition 4.2. The computation cost c is defined as the computation cost for training the model over the minibatch. This is approximately a constant times $(|\mathrm{S}_1|+|\mathrm{S}_2|)\cdot\bar{d}$, since we have to aggregate the neighbor information for nodes sampled in S_1 and S_2. We derive a lower bound on the pseudo precision rate $\rho(\xi)$ and observe that this bound converges to some value $\phi : (G, x_{v_1}, \ldots, x_{v_n}) \to \mathbb{R}$ which is independent of the batch size. Therefore, for any batch size $m < n$, there is some monotone decreasing function $\delta(m)$ such that $\rho(\xi) \geq \frac{1}{\phi(1+\delta)}$. The proof of the following proposition is provided in the supplementary material.

Proposition 4.1 *Let $\tilde{A} \in \mathbb{R}^{n\times n}$ be the normalized adjacency matrix of a graph $G = (V, E)$ with minimum degree $d_{\min} > \log n$, and suppose that $\max_{v,u\in V} |\tilde{A}_{u,v}| = O(1)$ and for each $v \in V$ the attribute $x_v = O(1)$. Let S_1 and S_2 be two random sets such that for every $i \in \{1, 2\}$, every $v \in V$ is picked to S_i with probability $m/n > \log n/d_{\min}$, so that $\mathbf{E}_{\mathrm{S}_i}[|\mathrm{S}_i|] = m$, and suppose that Assumption A.1 holds* [1].

Let ξ be the estimator from Definition 4.1, where W is some weight matrix with $\max_{v,u} |W_{v,u}| = O(1)$. Then there exists $\phi : (G, x_{v_1}, \ldots, x_{v_n}) \to \mathbb{R}$ and a monotone decreasing function

$$\delta(m) = \frac{2\bar{d}\sum_{(v,u)\in E} \frac{\tilde{A}_{v,u}^2 W_{v,u} x_u^2}{\alpha_v |N(v)|^2}}{m \cdot \phi}, \tag{8}$$

such that for every $m < n$ the pseudo precision of ξ is

$$\rho(\xi) \geq \left(\phi(1 + \delta(m))\right)^{-1}. \tag{9}$$

Remark 4.1 *Note that for simplicity of the presentation, we assumed that the layers are the same size and that all the attributes x_v are scalars. The bound can be generalized to any dimension by summing the variances of the individual coordinates of ξ_i.*

From the expression above, we can see that the bound on the pseudo precision converges to $1/\phi$ as $\delta(m)$ decreases, so that for any accuracy $\delta > 0$, there exists some m^* such that the pseudo precision is at least $1/(\phi(1 + \delta))$.

[1] This assumption is made to simplify the computations by assuming a weak correlation between the intersection of the neighborhoods of any two vertices, and although not true in general, since the graph is large and the degree of the nodes is small compared to n, we expect to have very little correlation between the sampled neighborhoods.

Table 1: Datasets statistics.

Dataset	Nodes	Edges	Avg deg	Features	Labels	Train/Val/Test Split
Pubmed	19,717	44,338	4.5	500	3	0.6/ 0.2 /0.2
Reddit	232,965	11,606,919	100	602	41	0.66/ 0.1 /0.24
Ogbn-arxiv	169,343	1,166,243	14	128	40	0.53/ 0.17 /0.3
Ogbn-products	2,449,029	61,859,140	50	100	47	0.08/ 0.02 /0.9

Table 2: Hyper-parameter setting and our suggested optimal batch size scale.

Dataset	Hidden size	Batch size to test	Optimal batch size scale
Pubmed	128	256, 512, 1k, 2k, 4k, 8k, full	$\sim$ 4k
Reddit	128	256, 1k, 4k, 8k, 16k, 32k	$\sim$ 2k
Ogbn-arxiv	256	256, 1k, 4k, 8k, 12k, 16k, 32k	$\sim$ 12k
Ogbn-product	256	256, 1k, 4k, 8k, 16k, 32k, 48k, 64k	$\sim$ 48k

Although the expression for $\delta(m)$ is hard to parse in general, we show that for d-regular graphs (i.e., all degrees are d), Lemma 4.1 yields a simplified bound on $\delta(m)$:

Corollary 4.2 *Let $G = (V, E)$ be a d-regular graph. There exists $\phi : (G, x_{v_1}, \ldots, x_{v_n}) \to \mathbb{R}$, such that for any $\delta > 0$ there exists $m(\delta) = O(n/d\delta)$ for which*

$$\rho(\xi) \geq \frac{1}{\phi(1+\delta)}.$$

For practical purposes, although the graphs we deal with are not d-regular, we propose to set the batch size to approximately $n/\bar{d}$, where $\bar{d}$ is the average degree of the graph. The intuition behind this guideline is that with this choice the bound on the pseudo precision rate reaches 1/2 of its maximum value. Beyond this setting, there are diminishing returns — the required compute time is increasing, but the variance has been reduced sufficiently so that additional decreases do not improve accuracy.

5 Experiments

5.1 Experiment Setting

We evaluate our theoretical findings regarding batch size selection using three state-of-the-art algorithms: ClusterGCN Chiang et al. (2019) (graph-wise sampling), FastGCN Chen et al. (2018) (layer-wise sampling) and GraphSAINT Zeng et al. (2020) (graph-wise sampling). We test each of the above algorithms against four public datasets: Pubmed, Reddit, Ogbn-arxiv and Ogbn-products Hu et al. (2020). The statistics of each dataset are shown in Table 1. Despite the different statistics among different datasets, we can evaluate the optimal scale of batch size from our theoretical result, which is shown in Table 2. For Pubmed dataset, we repartition the data with a train/validation/test split ratio of $6 : 2 : 2$. We keep the original partition for the remaining datasets. For all the tests, we use 3 layers of GCN. We set learning rate to 0.01. For Pubmed and Reddit, we run training for 100 epochs. For Ogbn-arxiv and Ogbn-products, we run training for 200 epochs. We use "node" sampler in GraphSAINT. The remaining settings are shown in Table 2.

For ClusterGCN and GraphSAINT, we conduct our experiments based on their published github repositories. For FastGCN, we implemented a version that can utilize GPU computation. All of our experiments are tested on a server equipped with a NVIDIA Tesla V100 GPU (32GB memory), and Intel Xeon Gold 6140 CPU (2.30GHz).

5.2 Numerical Results

Fig. 1 shows the relation between validation accuracy and training time with different batch size setting for various datasets with various algorithms. We also mark the positions that achieves best validation accuracy, 95% of best accuracy and 99% of best accuracy.

Table 3, Table 4, Table 5 and Table 6 shows the detailed results at the epoch with best validation accuracy for Pubmed, Reddit, Ogbn-arxiv and Ogbn-products respectively.

Table 3: Training performance on Pubmed.

	ClusterGCN			FastGCN			GraphSAINT		
Batch size	Val ACC	Test ACC	Train time (s)	Val ACC	Test ACC	Train time (s)	Val ACC	Test ACC	Train time (s)
256	0.8724	0.8625	19.307	0.8428	0.8435	15.675	0.8775	0.8737	28.922
512	0.8734	0.8595	6.815	0.8431	0.8437	7.502	0.8813	0.8644	5.337
1k	0.8696	0.8605	3.599	0.8451	0.8394	23.199	0.8811	0.8651	4.365
2k	0.8686	0.8658	2.681	0.8461	0.8427	53.841	0.8813	0.8704	6.738
4k	0.8693	0.8608	5.701	0.8461	0.8480	44.201	0.8811	0.8644	4.842
8k	0.8698	0.8600	6.202	0.8502	0.8484	28.152	0.8770	0.8755	3.355
Full	0.869	0.8572	15.247	0.8413	0.8467	12.046	-	-	- [2]

[2] The implementation from GraphSAINT does not have a plain way to run full batch setting.

Table 4: Training performance on Reddit.

	ClusterGCN			FastGCN			GraphSAINT		
Batch size	Val ACC	Test ACC	Train time (s)	Val ACC	Test ACC	Train time (s)	Val ACC	Test ACC	Train time (s)
256	0.9603	0.9594	288.206	0.9298	0.9311	812.742	0.9120	0.9112	882.175
1k	0.9599	0.9590	169.126	0.9345	0.9359	649.033	0.9486	0.9384	416.153
4k	0.9605	0.9606	73.502	0.9457	0.9474	675.869	0.9581	0.9567	461.416
8k	0.9616	0.9606	81.993	0.9501	0.9509	732.887	0.9615	0.9613	468.866
16k	0.9624	0.9603	130.732	0.9524	0.9538	321.390	0.9624	0.9595	440.772
32k	0.9631	0.9616	125.297	0.9543	0.9555	629.552	0.9647	0.9638	423.458

Table 5: Training performance on ogbn-arxiv.

	ClusterGCN			FastGCN			GraphSAINT		
Batch size	Val ACC	Test ACC	Train time (s)	Val ACC	Test ACC	Train time (s)	Val ACC	Test ACC	Train time (s)
256	0.6999	0.6953	97.5214	0.7163	0.6963	615.346	0.5860	0.5006	175.638
1k	0.6983	0.6940	55.088	0.7179	0.7086	337.386	0.6134	0.5058	47.173
4k	0.7027	0.7002	21.534	0.7216	0.7087	297.414	0.633	0.5335	48.700
8k	0.7050	0.6949	23.369	0.7262	0.7191	274.321	0.6447	0.5643	44.955
12k	0.7058	0.6957	23.689	0.7292	0.7209	198.988	0.6548	0.5924	46.396
16k	0.7089	0.6989	38.551	0.7311	0.7214	242.614	0.6566	0.6040	44.5272
32k	0.7117	0.7000	72.810	0.7317	0.7240	216.189	0.6669	0.6438	75.406

Table 6: Training performance on ogbn-products.

	ClusterGCN			FastGCN			GraphSAINT		
Batch size	Val ACC	Test ACC	Train time (s)	Val ACC	Test ACC	Train time (s)	Val ACC	Test ACC	Train time (s)
256	0.9088	0.7515	2110.088	0.8969	0.7990	3358.080	0.8737	0.7215	5483.181
1k	0.9115	0.75305	767.686	0.8987	0.7906	2604.168	0.8901	0.7340	3967.143
4k	0.9120	0.7492	351.658	0.8968	0.7897	2628.846	0.9032	0.7503	2511.132
8k	0.9121	0.7507	160.694	0.9031	0.8028	2455.458	0.9115	0.7681	1620.621
16k	0.9124	0.7502	229.460	0.9077	0.8030	2489.157	0.9180	0.7784	1452.206
32k	0.9130	0.7491	357.535	0.9138	0.8134	2164.151	0.9201	0.7939	966.082
48k	0.9110	0.7545	328.153	0.9158	0.8115	2545.868	0.9219	0.7917	1023.901
64k	0.9087	0.7517	390.819	0.9185	0.8062	2520.848	0.9219	0.8046	1024.634

Under the batch size of 256, a typical setting for conventional deep learning models, all the scenarios show a slow convergence rate and a bad validation/test accuracy. When the batch size increases, both the training efficiency and testing accuracy improve substantially. The reason is that the accuracy of neighbor aggregation rapidly increases when batch size is small. For FastGCN and GraphSAINT, testing accuracy will increase as we increase the batch size. However, there is some turning point in each dataset, beyond which the testing accuracy does not increase too much. This confirms our theoretical conclusion about diminishing returns. Reddit is a typical example where testing accuracy is boosted to batch size of 8k and beyond that value, it grows slowly. On the other hand, there is a sweet-spot for the fastest convergence rate, not necessary aligned with the turning point of testing accuracy. Above sweet-spot, variance of gradients in each minibatch does not decrease too much as batch size grows but computation cost keep increasing. Below the sweet-spot, rapid reduction of variance results a faster convergence. Our suggested scale of optimal batch size usually falls close to those two important points. For CluserGCN, the testing accuracy does not change too much, but the sweet point for convergence is close to our suggested optimal batch size too. In general, our guideline suggests some batch size that is much larger than conventional batch size setting like 512 and it is close to the spot where we can get a decent testing accuracy with efficient convergence rate.

Interestingly, on Reddit dataset Chen et al. (2018); Zou et al. (2019); Zeng et al. (2020); Chiang et al. (2019) report that graph-wise sampling (testing accuracy of $0.96+$) performs much better than the layer-wise sampling (testing accuracy of around 0.93). We found that the difference mainly comes from the fact that experiments in graph-wise sampling have a better batch size setting (8k in GraphSAINT) while the layer-wise experiment set a small batch size (400 in FastGCN). In our experi-

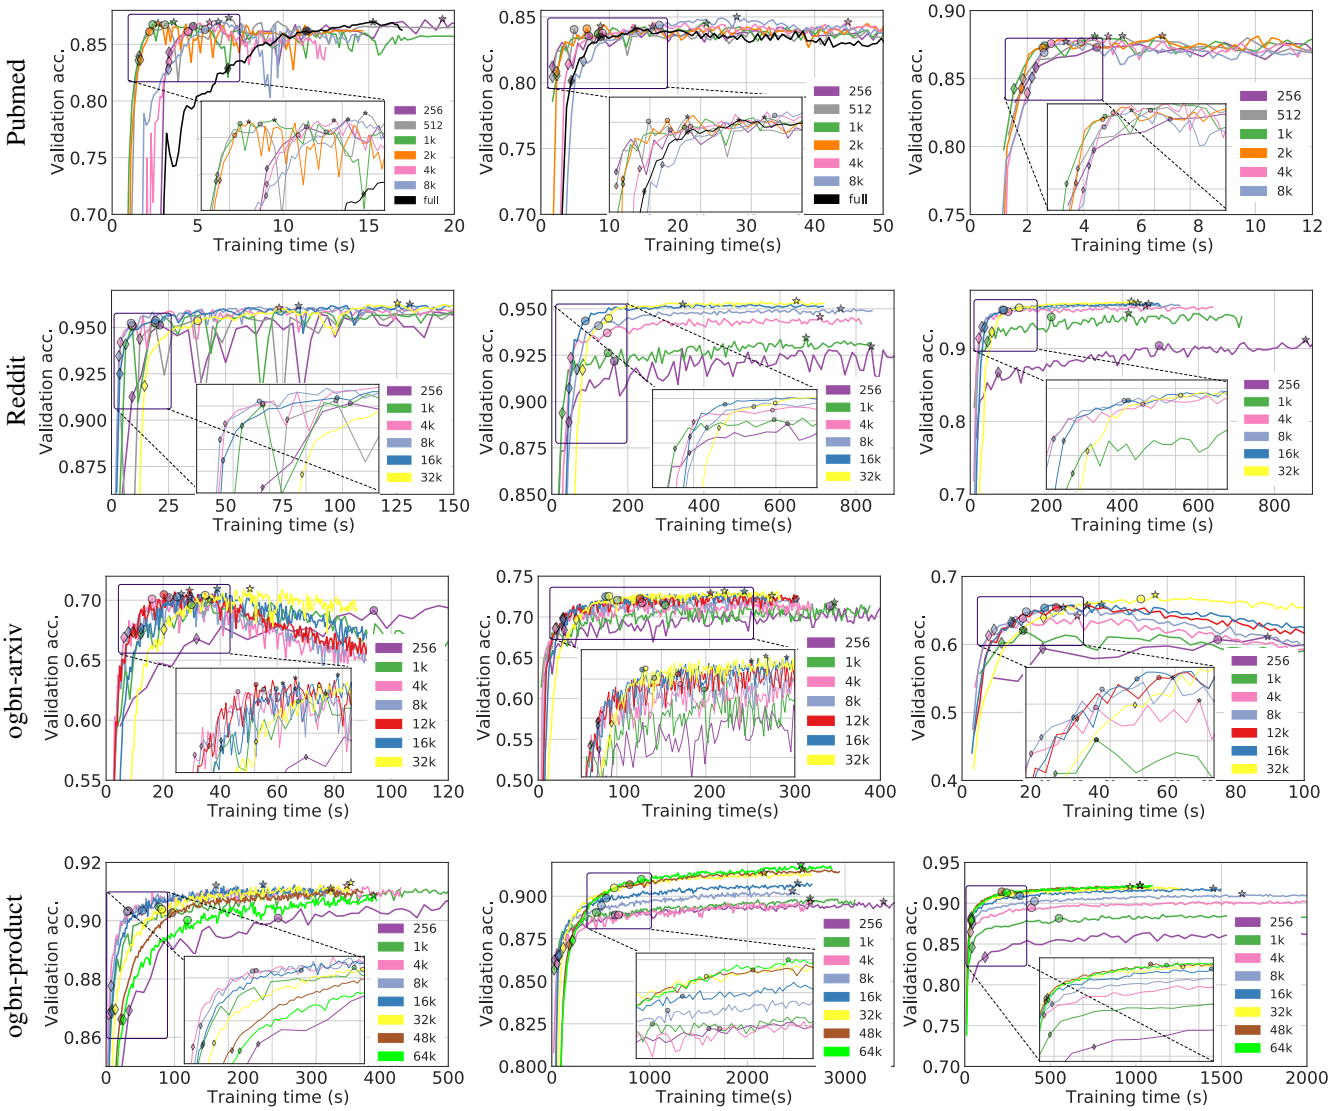

(a) ClusterGCN (b) FastGCN (c) GraphSAINT

Figure 1: Training time v.s. validation accuracy with different batch size settings for various datasets with various algorithms. $\star$: best validation ACC (BA), $\Diamond$: the first reach of 0.95 BA, $\circ$: the first reach of 0.99 BA.

ments, when we properly set the batch size, performances from layer-wise sampling and graph-wise sampling are close, which indicates the importance of batch size selection in GNN training.

6 Conclusion

We studied the batch size selection for SGD training of GNN models. We proposed *pseudo precision rate* metric that reflects training efficiency. We analyzed how the batch size influences this metric on an estimator that considers the randomness arising from two consecutive layers in GNN. By extensive experiments, we show that the batch size for GNN models should be much larger than typical setting of $\{4, 16, \ldots, 512\}$ from conventional deep learning models. With our suggested scale of batch size $n/\bar{d}$, n being the total number of nodes and $\bar{d}$ being the average node degree, GNN model can achieve decent testing performance efficiently.

REFERENCES

Motasem Alfarra, Slavomir Hanzely, Alyazeed Albasyoni, Bernard Ghanem, and Peter Richtarik. Adaptive learning of the optimal mini-batch size of sgd. *arXiv preprint arXiv:2005.01097*, 2020.

Jianfei Chen, Jun Zhu, and Le Song. Stochastic training of graph convolutional networks with variance reduction. In *International Conference on Machine Learning*, 2017.

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

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

Weilin Cong, Rana Forsati, Mahmut Kandemir, and Mehrdad Mahdavi. Minimal variance sampling with provable guarantees for fast training of graph neural networks. *arXiv preprint arXiv:2006.13866*, 2020.

Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *Advances in neural information processing systems*, pp. 3844–3852, 2016.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.

Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In *Proc. Int. Conf. Machine Learning*, 2017.

Robert Mansel Gower, Nicolas Loizou, Xun Qian, Alibek Sailanbayev, Egor Shulgin, and Peter Richtárik. Sgd: General analysis and improved rates. *arXiv preprint arXiv:1901.09401*, 2019.

Alex Graves. Generating sequences with recurrent neural networks. arxiv preprint arxiv: 1308.0850. 2013.

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

Fengxiang He, Tongliang Liu, and Dacheng Tao. Control batch size and learning rate to generalize well: Theoretical and empirical evidence. In *Advances in Neural Information Processing Systems*, pp. 1143–1152, 2019.

Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele Catasta, and Jure Leskovec. Open graph benchmark: Datasets for machine learning on graphs. *arXiv preprint arXiv:2005.00687*, 2020.

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

Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. On large-batch training for deep learning: Generalization gap and sharp minima. *arXiv preprint arXiv:1609.04836*, 2016.

Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.

Guohao Li, Matthias Muller, Ali Thabet, and Bernard Ghanem. Deepgcns: Can gcns go as deep as cnns? In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 9267–9276, 2019.

Guohao Li, Chenxin Xiong, Ali Thabet, and Bernard Ghanem. Deepergcn: All you need to train deeper gcns. In *International Conference on Machine Learning*, 2020.

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

Ziqi Liu, Chaochao Chen, Longfei Li, Jun Zhou, Xiaolong Li, Le Song, and Yuan Qi. Geniepath: Graph neural networks with adaptive receptive paths. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 4424–4431, 2019.

Dominic Masters and Carlo Luschi. Revisiting small batch training for deep neural networks. *arXiv preprint arXiv:1804.07612*, 2018.

Sam McCandlish, Jared Kaplan, Dario Amodei, and OpenAI Dota Team. An empirical model of large-batch training. *arXiv preprint arXiv:1812.06162*, 2018.

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

Art B. Owen. *Monte Carlo theory, methods and examples*. 2013.

Hongbin Pei, Bingzhe Wei, Kevin Chen-Chuan Chang, Yu Lei, and Bo Yang. Geom-gcn: Geometric graph convolutional networks. *arXiv preprint arXiv:2002.05287*, 2020.

Meng Qu, Yoshua Bengio, and Jian Tang. GMNN: Graph Markov neural networks. In *Proc. Int. Conf. Machine Learning*, 2019.

Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. Improving language understanding by generative pre-training, 2018.

Pavlo M Radiuk. Impact of training set batch size on the performance of convolutional neural networks for diverse datasets. *Information Technology and Management Science*, 20(1):20–24, 2017.

Leslie N Smith. A disciplined approach to neural network hyper-parameters: Part 1–learning rate, batch size, momentum, and weight decay. *arXiv preprint arXiv:1803.09820*, 2018.

Shikhar Vashishth, Prateek Yadav, Manik Bhandari, and Partha Talukdar. Confidence-based graph convolutional networks for semi-supervised learning. *arXiv preprint arXiv:1901.08255*, 2019.

Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017.

Rex Ying, Ruining He, Kaifeng Chen, Pong Eksombatchai, William L Hamilton, and Jure Leskovec. 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.

Hanqing Zeng, Hongkuan Zhou, Ajitesh Srivastava, Rajgopal Kannan, and Viktor Prasanna. Graph-SAINT: Graph sampling based inductive learning method. In *International Conference on Learning Representations*, 2020. URL `https://openreview.net/forum?id=BJe8pkHFwS`.

Yingxue Zhang, Soumyasundar Pal, Mark Coates, and Deniz Ustebay. Bayesian graph convolutional neural networks for semi-supervised classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 5829–5836, 2019.

Difan Zou, Ziniu Hu, Yewen Wang, Song Jiang, Yizhou Sun, and Quanquan Gu. Layer-dependent importance sampling for training deep and large graph convolutional networks. In *Advances in Neural Information Processing Systems*, pp. 11247–11256, 2019.

A BATCH SIZE ANALYSIS

In this section, we prove Lemma 4.1 and Corollary 4.2. We first present the following simplifying assumption that states that for any two vertices v_1, v_2 the intersection of their neighborhoods has weak correlation.

Assumption A.1 *Let $S \subset V$ be a random set drawn such that each $v \in V$ is picked to S with probability $p > \log n/d_{\min}$. There exist constants $c_1, c_2 \in \mathbb{R}$ such that*

$$\frac{1}{c_1}\mathop{\mathbf{E}}_{S}[|N(v_1)\cap S|]\mathop{\mathbf{E}}_{S}[|N(v_2)\cap S|] \le \mathop{\mathbf{E}}_{S}[|N(v_1)\cap S|\cdot|N(v_2)\cap S|] \le c_2\cdot\mathop{\mathbf{E}}_{S}[|N(v_1)\cap S|]\mathop{\mathbf{E}}_{S}[|N(v_2)\cap S|].$$

We start by stating a straightforward lemma from probability theory, whose proof is a direct application of Chernoff bounds.

Lemma A.1 *Let U be a finite set, and consider a random set $\mathrm{S} \subseteq U$ drawn such that each element in U is picked to S with probability p independently. Then, with probability at least $1-o(1/\mathrm{poly}(n))$*

$$|\mathrm{S}| \in \left[p|U| - \sqrt{p|U|\log n}, p|U| + \sqrt{p|U|\log n}\right]$$

For a fixed vertex u, we let $\tilde{x}_u = Wx_u$, so that our estimator

$$\xi = \frac{1}{|\mathrm{S}_1|}\sum_{v\in\mathrm{S}_1}\frac{\mathbf{1}_{N(v)\cap\mathrm{S}_2\neq\emptyset}}{|\mathrm{S}_2\cap N(v)|}\sum_{u\in N(v)\cap\mathrm{S}_2}\frac{\tilde{A}_{v,u}\cdot\tilde{x}_u}{\alpha_v}$$

We start with the proof of Proposition 4.1. Note that we considered the more general case where S_1 is picked according to probability $p = m_1/n$ and S_2 is picked according to probability $q = m_2/n$.

Proof of Proposition 4.1: We start with computing the mean of our estimator.

$$\begin{aligned}
\mathop{\mathbf{E}}_{\mathrm{S}_1,\mathrm{S}_2}\left[\frac{1}{|\mathrm{S}_1|}\sum_{v\in\mathrm{S}_1}\chi_v\right] &= \mathop{\mathbf{E}}_{\mathrm{S}_1,\mathrm{S}_2}\left[\frac{1}{|\mathrm{S}_1|}\sum_{v\in\mathrm{S}_1}\frac{\mathbf{1}_{N(v)\cap\mathrm{S}_2\neq\emptyset}}{\alpha_v\cdot|N(v)\cap\mathrm{S}_2|}\sum_{u\in N(v)\cap\mathrm{S}_2}\tilde{A}_{v,u}\tilde{x}_u\right] \\
&= \mathop{\mathbf{E}}_{\mathrm{S}_1,\mathrm{S}_2}\left[\frac{1}{|\mathrm{S}_1|}\sum_{v,u\in V}\mathbf{1}_{v\in\mathrm{S}_1}\cdot\mathbf{1}_{N(v)\cap\mathrm{S}_2\neq\emptyset}\cdot\mathbf{1}_{u\in\mathrm{S}_2\cap N(v)}\cdot\frac{\tilde{A}_{v,u}\tilde{x}_u}{\alpha_v\cdot|\mathrm{S}_2\cap N(v)|}\right] \\
&= \sum_{(v,u)\in E}\mathop{\mathbf{E}}_{\mathrm{S}_1,\mathrm{S}_2}\left[\frac{\mathbf{1}_{v\in\mathrm{S}_1}}{|\mathrm{S}_1|}\frac{\mathbf{1}_{N(v)\cap\mathrm{S}_2\neq\emptyset}\mathbf{1}_{u\in\mathrm{S}_2}}{\alpha_v\cdot|N(v)\cap\mathrm{S}_2|}\right]\cdot\tilde{A}_{v,u}\tilde{x}_u,
\end{aligned}$$

where for an event Z we let $\mathbf{1}_{\mathrm{Z}}$ denote the indicator random variable for Z.

Note that by the fact that each vertex is picked to S_1 (respectively S_2) independently w.p p we can apply Lemma A.1 and conclude that with very high probability $|\mathrm{S}_1| = pn \pm \sqrt{pn\log n} = \Theta(pn)$

By the independence of S_1 and S_2, and an application of Jensen's inequality, we can establish the following bound:

$$\mathop{\mathbf{E}}_{\mathrm{S}_1,\mathrm{S}_2}\left[\frac{\mathbf{1}_{v\in\mathrm{S}_1}}{|\mathrm{S}_1|}\frac{\mathbf{1}_{u\in\mathrm{S}_2}}{\alpha_v|N(v)\cap\mathrm{S}_2|}\right] = \mathop{\mathbf{E}}_{\mathrm{S}_1}\left[\frac{\mathbf{1}_{v\in\mathrm{S}_1}}{|\mathrm{S}_1|}\right]\mathop{\mathbf{E}}_{\mathrm{S}_2}\left[\frac{\mathbf{1}_{N(v)\cap\mathrm{S}_2\neq\emptyset}\mathbf{1}_{u\in\mathrm{S}_2}}{\alpha_v|N(v)\cap\mathrm{S}_2|}\right] \ge \Omega\left(\frac{\alpha_v}{n\alpha_v|N(v)|}\right) = \Omega\left(\frac{1}{n|N(v)|}\right),$$

where we used the fact that, $\mathbf{Pr}[\mathbf{1}_{N(v)\cap\mathrm{S}_2\neq\emptyset}] = \alpha_v$, leading to mean of $\Omega\left(\frac{1}{n}\sum_{(v,u)\in E}\frac{\tilde{A}_{v,u}\tilde{x}_u}{|N(v)|}\right)$.

Next, we compute the second moment of our estimator.

$$\begin{aligned}
&\mathop{\mathbf{E}}_{\mathrm{S}_1,\mathrm{S}_2}\left[\frac{1}{|\mathrm{S}_1|^2}\sum_{v_1,v_2\in\mathrm{S}_1}\chi_{v_1}\chi_{v_2}\right] \\
&= \mathop{\mathbf{E}}_{\mathrm{S}_1,\mathrm{S}_2}\left[\frac{1}{|\mathrm{S}_1|^2}\sum_{(v_1,u_1),(v_2,u_2)\in E}\frac{\mathbf{1}_{v_1,v_2\in\mathrm{S}_1}\cdot\mathbf{1}_{u_1,u_2\in\mathrm{S}_2}\cdot\mathbf{1}_{N(v_1)\cap\mathrm{S}_2\neq\emptyset}\mathbf{1}_{N(v_2)\cap\mathrm{S}_2\neq\emptyset}}{\alpha_{v_1}|N(v_1)\cap\mathrm{S}_2|\alpha_{v_2}|N(v_2)\cap\mathrm{S}_2|}\tilde{A}_{v_1,u_1}\tilde{x}_{u_1}\tilde{A}_{v_2,u_2}\tilde{x}_{u_2}\right] \\
&= \sum_{(v_1,u_1),(v_2,u_2)\in E}\mathop{\mathbf{E}}_{\mathrm{S}_1,\mathrm{S}_2}\left[\frac{\mathbf{1}_{v_1,v_2\in\mathrm{S}_1}\cdot\mathbf{1}_{u_1,u_2\in\mathrm{S}_2}\mathbf{1}_{N(v_1)\cap\mathrm{S}_2\neq\emptyset}\mathbf{1}_{N(v_2)\cap\mathrm{S}_2\neq\emptyset}}{|\mathrm{S}_1|^2\cdot\alpha_{v_1}\alpha_{v_2}|N(v_1)\cap\mathrm{S}_2||N(v_2)\cap\mathrm{S}_2|}\right]\cdot\tilde{A}_{v_1,u_1}\tilde{x}_{u_1}\tilde{A}_{v_2,u_2}\tilde{x}_{u_2}.
\end{aligned}$$

Similarly to before, we inspect the above expectation. In here, there are four cases corresponding to the following sets.

1. $\mathcal{C}_1 = \{(v_1,u_1),(v_2,u_2)\in E \mid v_1\neq v_2, u_1\neq u_2\}$: in a similar way to before, by Lemma A.1, we obtain

$$\begin{aligned}
&\mathop{\mathbf{E}}_{\mathrm{S}_1,\mathrm{S}_2}\left[\frac{\mathbf{1}_{v_1,v_2\in\mathrm{S}_1}\cdot\mathbf{1}_{u_1,u_2\in\mathrm{S}_2}\mathbf{1}_{N(v_1)\cap\mathrm{S}_2\neq\emptyset}\mathbf{1}_{N(v_2)\cap\mathrm{S}_2\neq\emptyset}}{|\mathrm{S}_1|^2\cdot\alpha_{v_1}\alpha_{v_2}|N(v_1)\cap\mathrm{S}_2||N(v_2)\cap\mathrm{S}_2|}\right] \\
&= \frac{p^2}{\alpha_{v_1}\alpha_{v_2}(pn)^2}\mathop{\mathbf{E}}_{\mathrm{S}_2}\left[\frac{\mathbf{1}_{N(v_1)\cap\mathrm{S}_2\neq\emptyset}\mathbf{1}_{N(v_2)\cap\mathrm{S}_2\neq\emptyset}\mathbf{1}_{u_1,u_2\in\mathrm{S}_2}}{|N(v_1)\cap\mathrm{S}_2||N(v_2)\cap\mathrm{S}_2|}\right] \\
&= \frac{p^2q^2}{(pn)^2}\mathop{\mathbf{E}}_{\mathrm{S}_2}\left[\frac{1}{|N(v_1)\cap\mathrm{S}_2||N(v_2)\cap\mathrm{S}_2|}\right],
\end{aligned}$$

   where we used the fact that for v_i, $\mathbf{Pr}[\mathbf{1}_{N(v_i)\cap\mathrm{S}_2\neq\emptyset}] = \alpha_{v_i}$.
   In order to analyze the above expectation, we use our assumption that neighborhoods are uncorrelated to get

$$\mathop{\mathbf{E}}_{\mathrm{S}_1,\mathrm{S}_2}\left[\frac{\mathbf{1}_{v_1,v_2\in\mathrm{S}_1}\cdot\mathbf{1}_{u_1,u_2\in\mathrm{S}_2}\mathbf{1}_{N(v_1)\cap\mathrm{S}_2\neq\emptyset}\mathbf{1}_{N(v_2)\cap\mathrm{S}_2\neq\emptyset}}{|\mathrm{S}_1|^2\cdot\alpha_{v_1}\alpha_{v_2}|N(v_1)\cap\mathrm{S}_2||N(v_2)\cap\mathrm{S}_2|}\right]\lesssim\frac{1}{n^2|N(v_1)|\cdot|N(v_2)|}.$$

2. $\mathcal{C}_2 = \{(v_1,u_1),(v_2,u_2)\in E \mid v_1\neq v_2, u_1 = u_2\}$: similarly to the previous case,

$$\mathop{\mathbf{E}}_{\mathrm{S}_1,\mathrm{S}_2}\left[\frac{\mathbf{1}_{v_1,v_2\in\mathrm{S}_1}\cdot\mathbf{1}_{u_1,u_2\in\mathrm{S}_2}\mathbf{1}_{N(v_1)\cap\mathrm{S}_2\neq\emptyset}\mathbf{1}_{N(v_2)\cap\mathrm{S}_2\neq\emptyset}}{|\mathrm{S}_1|^2\cdot\alpha_{v_1}\alpha_{v_2}|N(v_1)\cap\mathrm{S}_2||N(v_2)\cap\mathrm{S}_2|}\right]\lesssim\frac{1}{qn^2|N(v_1)||N(v_2)|}.$$

3. $\mathcal{C}_3 = \{(v_1,u_1),(v_2,u_2)\in E \mid v_1 = v_2, u_1\neq u_2\}$. This case requires extra care, since in this case, the neighborhoods of v_1 and v_2 are correlated (actually the same).

$$\mathop{\mathbf{E}}_{\mathrm{S}_1,\mathrm{S}_2}\left[\frac{\mathbf{1}_{v_1,v_2\in\mathrm{S}_1}\cdot\mathbf{1}_{u_1,u_2\in\mathrm{S}_2}\mathbf{1}_{N(v_1)\cap\mathrm{S}_2\neq\emptyset}\mathbf{1}_{N(v_2)\cap\mathrm{S}_2\neq\emptyset}}{|\mathrm{S}_1|^2\cdot\alpha_{v_1}\alpha_{v_2}|N(v_1)\cap\mathrm{S}_2||N(v_2)\cap\mathrm{S}_2|}\right]\simeq\frac{pq^2}{\alpha_v(pn)^2}\mathbf{E}\left[\frac{1}{|N(v)\cap\mathrm{S}_2|^2}\right].$$

   By Lemma A.1, with very high probability

$$|N(v)\cap\mathrm{S}_2|\in[q|N(v)|-\sqrt{q|N(v)|\log n}, q|N(v)|+\sqrt{q|N(v)|\log n}],$$

   and by our constraint that $q = \Omega(\log n/d_{\min}) = \Omega(\log n/|N(v)|)$, we have that with high probability

$$\mathop{\mathbf{E}}_{\mathrm{S}_2}\left[\frac{1}{|\mathrm{S}_2\cap N(v)|^2}\right]\simeq\frac{1}{(q|N(v)|\pm\sqrt{q|N(v)|\log n})^2}\lesssim O\left(\frac{1}{q^2|N(v)|^2}\right).$$

   so that

$$\frac{pq^2}{\alpha_v(pn)^2}\mathbf{E}\left[\frac{1}{|N(v)\cap\mathrm{S}_2|^2}\right]\lesssim\frac{1}{pn^2\alpha_v|N(v)|^2}$$

4. $\mathcal{C}_4 = \{(v_1,u_1),(v_2,u_2)\in E \mid v_1 = v_2, u_1 = u_2\}$. Similarly to the previous case,

$$\mathop{\mathbf{E}}_{\mathrm{S}_1,\mathrm{S}_2}\left[\frac{\mathbf{1}_{v_1,v_2\in\mathrm{S}_1}\cdot\mathbf{1}_{u_1,u_2\in\mathrm{S}_2}\mathbf{1}_{N(v_1)\cap\mathrm{S}_2\neq\emptyset}\mathbf{1}_{N(v_2)\cap\mathrm{S}_2\neq\emptyset}}{|\mathrm{S}_1|^2\cdot\alpha_{v_1}\alpha_{v_2}|N(v_1)\cap\mathrm{S}_2||N(v_2)\cap\mathrm{S}_2|}\right]=\frac{pq}{\alpha_v(pn)^2}\mathbf{E}\left[\frac{1}{|N(v)\cap\mathrm{S}|^2}\right]\lesssim\frac{1}{n^2\alpha_v pq|N(v)|^2}.$$

Combining the above and subtracting the expectation squared yields,

$$\mathrm{Var}_{\mathrm{S}_1,\mathrm{S}_2}\left[\frac{1}{|\mathrm{S}_1|}\sum_{v\in\mathrm{S}_1}\chi_v\right] \lesssim \frac{1}{n^2}\sum_{(v_1,u_1),(v_2,u_2)\in\mathcal{C}_2}\frac{\tilde{A}_{v_1,u_1}\tilde{A}_{v_2,u_1}\tilde{x}_{u_1}^2}{q|N(v_1)||N(v_2)|}$$
$$+\frac{1}{n^2}\sum_{(v_1,u_1),(v_2,u_2)\in\mathcal{C}_3}\frac{\tilde{A}_{v_1,u_1}\tilde{x}_{u_1}\tilde{A}_{v_1,u_2}\tilde{x}_{u_2}}{p\alpha_v|N(v)|^2}+\frac{1}{n^2}\sum_{(v_1,u_1),(v_2,u_2)\in\mathcal{C}_4}\frac{\tilde{A}_{v,u}^2x_{u_1}^2}{pq\alpha_v|N(v)|^2}-\left(\frac{1}{n}\sum_{(v,u)\in E}\frac{\tilde{A}_{v,u}\tilde{x}_u}{|N(v)|}\right)^2$$

After rearranging we get

$$\frac{1}{n^2}\Bigg(\sum_{(v_1,u_1),(v_2,u_2)\in\mathcal{C}_2}\left(\frac{1}{q|N(v_1)||N(v_2)|}-\frac{1}{|N(v_1)||N(v_2)|}\right)\tilde{A}_{v_1,u_1}\tilde{A}_{v_2,u_1}\tilde{x}_{u_1}^2$$
$$+\sum_{(v_1,u_1),(v_2,u_2)\in\mathcal{C}_3}\left(\frac{1}{p\alpha_v|N(v)|^2}-\frac{1}{|N(v)|^2}\right)\tilde{A}_{v_1,u_1}\tilde{x}_{u_1}\tilde{A}_{v_1,u_2}\tilde{x}_{u_2}+\sum_{(v,u)\in E}\left(\frac{1}{pq\alpha_v|N(v)|^2}-\frac{1}{|N(v)|^2}\right)\tilde{A}_{v,u}^2\tilde{x}_u^2\Bigg).$$

If we let $p=m_1/n$ and $q=m_2/n$ so that $\mathbf{E}_{\mathrm{S}_1}[|\mathrm{S}_1|]=m_1$ and $\mathbf{E}_{\mathrm{S}_2}[|\mathrm{S}_2|]=m_2$, we get that

$$\mathrm{Var}_{\mathrm{S}_1,\mathrm{S}_2}\left[\frac{1}{|\mathrm{S}_1|}\sum_{v\in\mathrm{S}_1}\chi_v\right]$$
$$\lesssim\frac{1}{n^2}\Bigg(\sum_{(v_1,u_1),(v_2,u_2)\in\mathcal{C}_2}\frac{\tilde{A}_{v_1,u_1}\tilde{A}_{v_2,u_1}\tilde{x}_{u_1}^2}{q|N(v_1)||N(v_2)|}+\sum_{(v_1,u_1),(v_2,u_2)\in\mathcal{C}_3}\frac{\tilde{A}_{v_1,u_1}\tilde{x}_{u_1}\tilde{A}_{v_1,u_2}\tilde{x}_{u_2}}{p\alpha_v|N(v)|^2}+\sum_{(v,u)\in E}\frac{\tilde{A}_{v,u}^2\tilde{x}_u^2}{pq\alpha_v|N(v)|^2}\Bigg)$$
$$\lesssim\sum_{(v_1,u_1),(v_2,u_2)\in\mathcal{C}_2}\left(\frac{1}{nm_2}\right)\frac{\tilde{A}_{v_1,u_1}\tilde{A}_{v_2,u_1}\tilde{x}_{u_1}^2}{|N(v_1)||N(v_2)|}+\sum_{(v_1,u_1),(v_2,u_2)\in\mathcal{C}_3}\left(\frac{1}{m_1\cdot n}\right)\frac{\tilde{A}_{v_1,u_1}\tilde{x}_{u_1}\tilde{A}_{v_1,u_2}\tilde{x}_{u_2}}{\alpha_v|N(v)|^2}$$
$$+\sum_{(v,u)\in E}\left(\frac{1}{m_1\cdot m_2}\right)\frac{\tilde{A}_{v,u}^2\tilde{x}_u^2}{\alpha_v|N(v)|^2}.$$

Note that the variance decreases as m_1, m_2 increase. Therefore, if we assume for simplicity that $m_1=m_2=m$ the variance is bounded by

$$\mathrm{Var}_{\mathrm{S}_1,\mathrm{S}_2}[\boldsymbol{\xi}]\lesssim\frac{1}{nm}\Bigg(\sum_{(v_1,u_1),(v_2,u_2)\in\mathcal{C}_2}\frac{\tilde{A}_{v_1,u_1}\tilde{A}_{v_2,u_1}\tilde{x}_{u_1}^2}{|N(v_1)||N(v_2)|}+\sum_{(v_1,u_1),(v_2,u_2)\in\mathcal{C}_3}\frac{\tilde{A}_{v_1,u_1}\tilde{x}_{u_1}\tilde{A}_{v_1,u_2}\tilde{x}_{u_2}}{\alpha_v|N(v)|^2}\Bigg)$$
$$+\sum_{(v,u)\in E}\left(\frac{1}{m^2}\right)\frac{\tilde{A}_{v,u}^2\tilde{x}_u^2}{\alpha_v|N(v)|^2}. \tag{10}$$

Let's consider the pseudo precision of our estimator. Note that the cost generating the estimator ξ is approximately $2m\cdot\bar{d}$, where $\bar{d}$ is the average degree of the graph. This is since we have roughly $2m$ vertices to generate and for each vertex sampled we need to aggregate its neighbors information (which is approximately $\bar{d}$).

$$\rho(\xi)=(\mathrm{Var}(\xi)\cdot\mathrm{Cost}(\xi))^{-1}\gtrsim\Bigg(\frac{2\bar{d}}{n}\Bigg(\sum_{(v_1,u_1),(v_2,u_2)\in\mathcal{C}_2}\frac{\tilde{A}_{v_1,u_1}\tilde{A}_{v_2,u_1}\tilde{x}_{u_1}^2}{|N(v_1)||N(v_2)|}+\sum_{(v_1,u_1),(v_2,u_2)\in\mathcal{C}_3}\frac{\tilde{A}_{v_1,u_1}\tilde{x}_{u_1}\tilde{A}_{v_1,u_2}\tilde{x}_{u_2}}{\alpha_v|N(v)|^2}\Bigg)$$
$$+\sum_{(v,u)\in E}\left(\frac{2\bar{d}}{m}\right)\frac{\tilde{A}_{v,u}^2\tilde{x}_u^2}{\alpha_v|N(v)|^2}\Bigg)^{-1}. \tag{11}$$

Let

$$\phi = \frac{2\bar{d}}{n}\left(\sum_{(v_1,u_1),(v_2,u_2)\in\mathcal{C}_2} \frac{\tilde{A}_{v_1,u_1}\tilde{A}_{v_2,u_1}\tilde{x}_{u_1}^2}{|N(v_1)||N(v_2)|} + \sum_{(v_1,u_1),(v_2,u_2)\in\mathcal{C}_3} \frac{\tilde{A}_{v_1,u_1}\tilde{x}_{u_1}\tilde{A}_{v_1,u_2}\tilde{x}_{u_2}}{\alpha_v|N(v)|^2}\right),$$

and note that as m increase, the efficiency converges to ϕ.

If we define δ as

$$\delta(m) = \frac{2\bar{d}\sum_{(v,u)\in E}\frac{\tilde{A}_{v,u}^2\tilde{x}_u^2}{\alpha_v|N(v)|^2}}{m\phi}, \tag{12}$$

we get that $\rho(\xi) \geq \frac{1}{\phi(1+\delta)}$, as claimed. ■

Now we show that if we assume that the graph is d-regular, we can get a clean relation between the efficiency of our estimator and the size of the batch.

Proof of Corollary 4.2: Fix any $\delta > 0$. By Equation (12),

$$m = \frac{2\bar{d}\sum_{(v,u)\in E}\frac{\tilde{A}_{v,u}^2\tilde{x}_u^2}{\alpha_v|N(v)|^2}}{\delta\phi}$$

so that the pseudo precision is at least

$$\rho(\xi) \geq \left(\phi(1+\delta(m))\right)^{-1}.$$

By assuming that the graph is d-regular, and using the fact that for all $v \in V$, $1 - e^{-qd} \leq \alpha_v \leq qd$ and $qd > 1$,

$$m = O\left(\frac{1}{\delta}\cdot\frac{n\cdot\left(\frac{1}{1-e^{-qd}}\frac{nd}{d^2}\right)}{\frac{1}{d^2}n\cdot\binom{d}{2}+\frac{1}{dq}\frac{1}{d^2}n\binom{d}{2}}\right) = O\left(\frac{1}{\delta}\cdot\frac{n/d\cdot\frac{1}{1-e^{-qd}}}{\frac{1}{2}+\frac{1}{2dq}}\right) = O\left(\frac{n}{\delta d}\right).$$

■

Remark on graph-wise sampling: Note that the case where $\mathrm{S}_1 = \mathrm{S}_2$ is equivalent to subgraph sampling. The same analysis can be adapted to subgraph sampling, and yields the same proposed minibatch size.