(LA)YER-NEIGH(BOR) SAMPLING: DEFUSING NEIGHBORHOOD EXPLOSION

Anonymous authors
Paper under double-blind review

ABSTRACT

Graph Neural Networks have recently received a significant attention, however, training them at a large scale still remains as a challenge. Minibatch training coupled with sampling is used to alleviate this challenge. However existing approaches either suffer from the neighborhood explosion phenomenon or does not have good performance. To deal with these issues, we propose a new sampling algorithm called LAyer-neighBOR sampling (LABOR). It is designed to be a direct replacement for Neighborhood Sampling with the same fanout hyperparameter while sampling much fewer vertices, without sacrificing quality. By design, the variance of the estimator of each vertex matches Neighbor Sampling from the point of view from a single vertex. In our experiments, we demonstrate the superiority of our approach when it comes to model convergence behaviour against Neighbor Sampling and also the other Layer Sampling approaches under the same limited vertex sampling budget constraints.

1 INTRODUCTION

Graph Neural Networks (GNN) Hamilton et al. (2017); Kipf & Welling (2017) have become de facto models for representation learning on graph structured data. Hence they have started being deployed in production systems Ying et al. (2018); Niu et al. (2020). These models iteratively update the node embeddings by passing messages along the direction of the edges in the given graph with nonlinearities in between different layers. With l layers, the computed node embedding contain information from the l-hop neighborhood of the seed vertex.

In the production setting, the GNN models need to be trained on billion-scale graphs (Ching et al., 2015; Ying et al., 2018). The training of these models takes hours to days even on distributed systems Zheng et al. (2022b;a). As in general Deep Neural Networks (DNN), it is more efficient to use minibatch training (Bertsekas, 1994) on GNNs, even though it is a bit trickier in this case. The node embeddings in GNNs depend recursively on their set of neighbors' embeddings, so when there are l layers, this dependency spans the l-hop neighborhood of the node. Real world graphs usually have a very small diameter and if l is large, the l-hop neighborhood may very well span the entire graph, also known as the Neighborhood Explosion Phenomenon (NEP) (Zeng et al., 2020).

To solve these issues, researchers proposed sampling a subgraph of the l-hop neighborhood of the nodes in the batch. There are mainly three different approaches: Node-based, Layer-based and Subgraph-based methods. Node-based sampling methods (Hamilton et al., 2017; Chen et al., 2018; Liu et al., 2020; Zhang et al., 2021) sample independently and recursively for each node. It was noticed that node-based methods sample subgraphs that are too shallow, i.e., with a low ratio of number of edges to nodes. Thus layer-based sampling methods were proposed (Chen et al., 2018; Zou et al., 2019; Huang et al., 2018; Dong et al., 2021), where the sampling for the whole layer is done collectively. On the other hand subgraph sampling methods (Chiang et al., 2019; Zeng et al., 2020; Hu et al., 2020) do not use the recursive layer by layer sampling scheme used in the node- and layer-based sampling methods and instead tend to use the same subgraph for all of the layers. Some of these sampling methods take the magnitudes of embeddings into account (Liu et al., 2020; Zhang et al., 2021; Huang et al., 2018), while others, such as Chen et al. (2018); Cong et al. (2021), cache the historical embeddings to reduce the variance of the computed approximate embeddings. There are methods sampling from a vertex cache Dong et al. (2021) filled with popular vertices. Most

of these approaches are orthogonal to each other and they can be incorporated into other sampling algorithms.

Node-based sampling methods suffer most from the NEP but they guarantee a good approximation for each embedding by ensuring each vertex gets k neighbors which is the only hyperparameter of the sampling algorithm. Layer-based sampling methods do not suffer as much from the NEP because number of vertices sampled is a hyperparameter but they can not guarantee that each vertex approximation is good enough and also their hyperparameters are hard to reason with, number of nodes to sample at each layer depends highly on the graph structure (as the numbers in Table 2 show). Subgraph sampling methods usually have more bias than their node- and layer-based counterparts. Hence, in this paper, we focus on the node- and layer-based sampling methods and combine their advantages. The major contributions of this work can be listed as follows:

- We propose a new sampling algorithm called LABOR, combining advantages of neighbor and layer sampling approaches using Poisson Sampling. LABOR correlates the sampling procedures of the given set of seed nodes so that the sampled vertices from different seeds have a lot of overlap, resulting into a big reduction in computation, memory and communication. Furthermore, LABOR has the same hyperparameters as neighbor sampling to use as a drop-in replacement.
- We experimentally verify our findings, show that our proposed sampling algorithm LABOR outperforms both neighbor sampling and layer sampling approaches.

2 Background

Graph Neural Networks: Given a directed graph $\mathcal{G} = (V, E)$, where V and $E \subset V \times V$ are vertex and edge sets respectively, $(t \to s) \in E$ denotes an edge from a source vertex $t \in V$ to a destination vertex $s \in V$, and A_{ts} denotes the corresponding edge weight if provided. If we have a batch of seed vertices $S \subset V$, let us define l-hop neighborhood $N^l(S)$ for the incoming edges as follows:

$$N(s) = \{t | (t \to s) \in E\},\ N^1(S) = N(S) = \cup_{s \in S} N(s),\ N^l(S) = N(N^{l-1}(S)) \tag{1}$$

Let us also define the degree d_s of vertex s as $d_s = |N(s)|$. To simplify the discussion, let's assume uniform edge weights, $A_{ts} = 1, \forall (t \to s) \in E$. Then, our goal is to estimate the following for each vertex $s \in S$, where $H_t^{(l-1)}$ is defined as the embedding of the vertex t at layer $l-1$, and $W^{(l-1)}$ is the trainable weight matrix at layer $l-1$, and σ is the nonlinear activation function (Hamilton et al., 2017):

$$Z_s^{(l)} = \frac{1}{d_s} \sum_{t \to s} H_t^{(l-1)} W^{(l-1)},\ H_s^{(l)} = \sigma(Z_s^{(l)}) \tag{2}$$

Exact Stochastic Gradient Descent: If we have a node prediction task and $V_t \subseteq V$ is the set of training vertices, $y_s, s \in V_t$ are the labels of the prediction task, and ℓ is the loss function for the prediction task, then our goal is to minimize the following loss function: $\frac{1}{|V_t|} \sum_{s \in V_t} \ell(y_s, Z_s^l)$. Replacing V_t in the loss function with $S \subset V_t$ for each iteration of gradient descent, we get stochastic gradient descent for GNNs. However with l layers, the computation dependency is on $N^l(S)$, which reaches large portion of the real world graphs, i.e. $|N^l(S)| \approx |V|$, making each iteration costly both in terms of computation and memory.

Neighbor Sampling: Neighbor sampling approach was proposed by Hamilton et al. (2017) to approximate $Z_s^{(l)}$ for each $s \in S$ with a subset of $N^l(S)$. Given a fanout hyperparameter k, this subset is computed recursively by randomly picking k neighbors for each $s \in S$ from $N(s)$ to form the next layer S^1, that is a subset of $N^1(S)$. If $d_s \leq k$, then the exact neighborhood $N(s)$ is used. For the next layer, S^1 is treated as the new set of seed vertices and this procedure is applied recursively.

Revisiting LADIES, Dependent Layer-based Sampling From now on, we will drop the layer notation and focus on a single layer and also ignore the nonlinearities. Let us define $M_t = H_t W$ as a shorthand notation. Then our goal is to approximate:

$$H_s = \frac{1}{d_s} \sum_{t \to s} M_t \tag{3}$$

If we assign probabilities $\pi_t > 0, \forall t \in N(S)$ and normalize it so that $\sum_{t \in N(S)} \pi_t = 1$, then use sampling with replacement to sample $T \subset N(S)$ with $|T| = n$, where n is the number of vertices to sample given as input to the LADIES algorithm and T is a multiset possibly with multiple copies of the same vertices, and let $\tilde{d}_s = |T \cap N(s)|$ which is the number of sampled vertices for a given vertex s, we get the following two possible estimators for each vertex $s \in S$:

$$H'_s = \frac{1}{n d_s} \sum_{t \in T \cap N(s)} \frac{M_t}{\pi_t} \tag{4}$$

$$H''_s = \frac{\sum_{t \in T \cap N(s)} \frac{M_t}{\pi_t}}{\sum_{t \in T \cap N(s)} \frac{1}{\pi_t}} \tag{5}$$

Note that H'_s in Eq. 4 is the Thompson-Horvitz estimator and the H''_s in Eq. 5 is the Hajek estimator. For a comparison between the two and how to get an even better estimator by combining them, see Khan & Ugander (2021). The formulation in the LADIES paper uses H'_s, but it proposes to row-normalize the sampled adjacency matrix, meaning they use H''_s in their implementation. However, analysing the variance of the Thompson-Horvitz estimator is simpler and its variance serves as an upper bound for the variance of the Hajek estimator when $|M_t|$ and π_t are uncorrelated Khan & Ugander (2021); Dorfman (1997), which we assume to be true in our case.

$$\mathrm{Var}(H''_s) \leq \mathrm{Var}(H'_s) = \frac{1}{\tilde{d}_s d_s^2} \sum_{t \to s} \pi_t \sum_{t' \to s} \frac{\mathrm{Var}(M_{t'})}{\pi_{t'}} \tag{6}$$

Since we assume that we do not have access to the computed embeddings and to simplify the analysis, we assume that $\mathrm{Var}(M_t) = 1$ from now on. One can see that $\mathrm{Var}(H'_s)$ is minimized when $\pi_t = p, \forall t \to s$ under the constraint $\sum_{t \to s} \pi_t \leq p d_s$ for some constant $p \in [0, 1]$, hence any deviation from uniformity increases the variance. The variance is also smaller the larger $\tilde{d}_s$ is. However, in theory and in practice, there is no guarantee that each vertex $s \in S$ will get any neighbors in T, not to mention equal numbers of neighbors. Some vertices will have pretty good estimators with thousands of samples and very low variances, while others might not even get a single neighbor sampled. For this reason, we designed LABOR so that every vertex in S will sample enough neighbors in expectation.

While LADIES is optimal from an approximate matrix multiplication perspective Chen et al. (2022), it is far from optimal in the case of nonlinearities and multiple layers. If there was a single layer, then the loss is computed by summing up the embeddings of the seed vertices, and it is optimal in this case. However, the existence of nonlinearities in-between layers and the fact that there are multiple layers necessitates that each vertex gets a good enough estimator with low enough variance. Also, LADIES gives a formulation using sampling with replacement instead of without replacement and that is sub-optimal from the variance perspective while its implementation uses sampling without replacement without taking care of the bias created thereby. In the next section, we will show how all of these problems are addressed by our newly proposed poisson sampling framework and LABOR sampling.

3 LABOR: LAYER NEIGHBOR SAMPLING

As mentioned previously, node-based sampling methods suffer from sampling too shallow subgraphs leading to NEP in just a few hops (e.g., see Table 2). Layer sampling methods Zou et al. (2019) attempt to fix this by sampling a fixed number of vertices in each layer, however they can not ensure that the estimators for the vertices are of high quality, and it is hard to reason how to choose the number of vertices to sample in each layer. The original paper of LADIES Zou et al. (2019) proposes using the same number of vertices in each layer while papers evaluating it found it is better to sample an increasing number of vertices in each layer Liu et al. (2020); Chen et al. (2022). There

is no systematic way to choose how many vertices to sample in each layer for the LADIES method, and since each graph has different density and connectivity structure, this choice highly depends on the graph in question. Therefore, due to its simplicity and high quality results, Neighbor Sampling currently seems to be the most popular of all of the sampling approaches and there exists high quality implementations on both CPUs and GPUs in the popular GNN frameworks Wang et al. (2019); Fey & Lenssen (2019).

We propose a new approach that combines the advantages of layer and neighbor sampling approaches using a vertex-centric variance based framework, reducing the number of sampled vertices drastically while ensuring the training quality does not suffer and matches the quality of neighbor sampling. Another advantage of our method is that the user only needs to choose the batch size and the fanout hyperparameters as in the Neighbor Sampling approach, the algorithm itself then samples the minimum number of vertices in the later layers in an unbiased way while ensuring each vertex gets enough neighbors and a good approximation.

3.1 LABOR Sampling

The design philosophy of LABOR Sampling is to create a direct alternative to Neighbor Sampling while incorporating the advantages of layer sampling. In layer sampling, the main idea can be summarized as individual vertices making correlated decisions while sampling their neighbors, because in the end if a vertex t is sampled, all edges $t \to s$, into the seed vertices S ($s \in S$) are added to the sampled subgraph. This can be interpreted as vertices in S making a collective decision on whether to sample t.

The other thing to keep in mind is that, the existing layer sampling methods use sampling with replacement when doing importance sampling with unequal probabilities, because it is computationally intractable to compute the inclusion probabilities in the without replacement case. The Hajek estimator in the without replacement case with equal probabilities becomes:

$$H_s'' = \frac{\sum_{t \in T \cap N(s)} \frac{M_t}{\bar{\pi}_t}}{\sum_{t \in T \cap N(s)} \frac{1}{\bar{\pi}_t}} = \frac{\sum_{t \in T \cap N(s)} M_t |N(S)|}{\sum_{t \in T \cap N(s)} |N(S)|} = \frac{1}{\tilde{d}_s} \sum_{t \in T \cap N(s)} M_t \tag{7}$$

and it has the variance:

$$\mathrm{Var}(H_s'') = \frac{d_s - \tilde{d}_s}{d_s - 1} \frac{1}{\tilde{d}_s} \tag{8}$$

Keeping these two points in mind, we use Poisson Sampling and design LABOR sampling around it. First, let us show how one can do layer sampling using Poisson sampling (PLADIES). Given probabilities $\pi_t \in [0, 1], \forall t \in N(S)$ so that $\sum_{t \in N(S)} \pi_t = n$, we include $t \in N(S)$ in our sample T with probability π_t by flipping a coin for it, i.e., we sample $r_t \sim U(0, 1)$ and include $t \in T$ if $r_t \leq \pi_t$. In the end, $E[|T|] = n$ and we can still use the Hajek estimator H_s'' or the Horvitz Thomson estimator H_s' to estimate H_s. This way of doing layer sampling is unbiased by construction and achieves the same goal in linear time in constrast to the quadratic time debiasing approach explained in Chen et al. (2022). In this case, the variance becomes (Williams et al., 1998):

$$\mathrm{Var}(H_s'') \leq \mathrm{Var}(H_s') = \frac{1}{d_s^2} \sum_{t \to s} \frac{1}{\pi_t} - \frac{1}{d_s} \tag{9}$$

One can notice the existence of the minus term $\frac{1}{d_s}$, and it enables the variance to converge to 0 if all $\pi_t = 1$ and we get the exact result. However, in the sampling with replacement case, the variance goes to 0 only as the sample size goes to infinity.

This type of mimicking Layer Sampling with Poisson Sampling still has the disadvantage that $\tilde{d}_s$ varies wildly for different s. To overcome this and mimic Neighbor Sampling where $E[\tilde{d}_s] = \min(d_s, k)$, where k is a given fanout hyperparameter. For given $\pi_t \geq 0, \forall t \in N(S)$ denoting unnormalized probabilities, for a given s, let us define c_s as the quantity satisfying the following equality if $k < d_s$, otherwise $c_s = \max_{t \to s} \frac{1}{\pi_t}$:

$$\frac{1}{d_s^2} \sum_{t \to s} \frac{1}{\min(1, c_s \pi_t)} - \frac{1}{d_s} = \frac{1}{k} - \frac{1}{d_s} \tag{10}$$

Note that $\frac{1}{k} - \frac{1}{d_s}$ is the variance when $\pi_t = \frac{k}{d_s}, \forall t \in N(s)$ so that $E[\tilde{d}_s] = k$. Also note that:

$$\frac{1}{k} - \frac{1}{d_s} - \frac{d_s - k}{d_s - 1}\frac{1}{k} = \frac{d_s - k}{kd_s} - \frac{d_s - k}{d_s - 1}\frac{1}{k} = \frac{d_s - k}{k}(\frac{1}{d_s} - \frac{1}{d_s - 1}) < 0 \tag{11}$$

meaning that the variance target we set through Eq. 10 is strictly better than Neighbor Sampling's variance in Eq. 8 and it will result in $E[\tilde{d}_s] \geq k$ with strict equality in the uniform probability case. Then each vertex $s \in S$ samples $t \to s$ with probability $c_s\pi_t$. To keep the collective decision making, we sample $r_t \sim U(0, 1), \forall t \in N(S)$ and vertex s samples vertex t if and only if $r_t \leq c_s\pi_t$. Note that if we use a uniform random variable for each edge r_{ts} instead of each vertex r_t, and if π is uniformly initialized, then we get the same behaviour as Neighbor Sampling.

3.2 Importance Sampling

Given the sampling procedure above, one wonders how different choices of $\pi \geq 0$ will affect $|T|$, the total number of unique vertices sampled. In our case, it is extremely easy to compute:

$$E[|T|] = \sum_{t \in N(S)} \mathbb{P}(t \in T) = \sum_{t \in N(S)} \min(1, \pi_t \max_{t \to s} c_s) \tag{12}$$

In particular, we need to find $\pi^* \geq 0$ minimizing $E[|T|]$:

$$\pi^* = \arg\min_{\pi \geq 0} \sum_{t \in N(S)} \min(1, \pi_t \max_{t \to s} c_s) \tag{13}$$

Note that for any given $\pi \geq 0$, $E[|T|]$ is the same for any vector multiple $x\pi, x \in \mathbb{R}^+$, meaning that the objective function is homogenous of degree 0.

3.3 Computing c and π^*

Note that c_s was defined to be the scalar satisfying the following equality involving the variance of the estimator of H_s:

$$\frac{1}{d_s^2} \sum_{t \to s} \frac{1}{\min(1, c_s\pi_t)} - \frac{1}{d_s} = \frac{1}{k} - \frac{1}{d_s} \tag{14}$$

If we rearrange the terms, we get:

$$\sum_{t \to s} \frac{1}{\min(1, c_s\pi_t)} = \frac{d_s^2}{k} \tag{15}$$

One can see that the left hand side of the equality is monotonically decreasing with respect to $c_s \geq 0$. Thus one can use binary search to find the c_s satisfying the above equality to any precision needed. But we opt to use the following iterative algorithm to compute it:

$$v_s^{(0)} = 0, c_s^{(0)} = \frac{k}{d_s^2} \sum_{t \to s} \frac{1}{\pi_t} \tag{16}$$

$$c_s^{(i+1)} = \frac{c_s^{(i)}}{\frac{d_s^2}{k} - v_s^{(i)}} \Big(- v_s^{(i)} + \sum_{t \to s} \frac{1}{\min(1, c_s^{(i)}\pi_t)} \Big), v_s^{(i+1)} = \sum_{t \to s} \mathbb{1}[c_s^{(i+1)}\pi_t \geq 1] \tag{17}$$

This iterative algorithm converges in at most d_s steps and the convergence is exact and monotonic from below. One can also implement it in linear time $\mathcal{O}(d_s)$ if $\{\pi_t \mid t \to s\}$ is sorted and making use of precomputed prefix sum arrays. Note that $c = c(\pi)$, meaning that c is a function of the given probability vector π. To compute π^*, we use a similar fixed point iteration as follows:

$$\pi^{(0)} = 1, \forall t \in N(S) : \pi_t^{(i+1)} = \pi_t^{(i)} \max_{t \to s} c_s(\pi^{(i)}) \tag{18}$$

Thus, we alternate between computing $c = c(\pi)$, meaning c is computed with the current π, and updating the π with the computed c values. Each step of this iteration is guaranteed to lower the objective function value in Eq. 13 until convergence to a fixed point, see the Appendix A.1. Modified method for a given nonuniform weight matrix A_{ts} is explained in the Appendix A.3.

Table 1: Datasets used in experiments, numbers of vertices, edges, avg. degree, features, sampling budget used, training, validation and test vertex split.

Dataset	$\lvert V\rvert$	$\lvert E\rvert$	$\frac{\lvert E\rvert}{\lvert V\rvert}$	**# feats.**	$\lvert V_3\rvert$ **budget**	**train - val - test (%)**
products	2.45M	61.9M	25.26	100	60k	8 - 2 - 90
reddit	233K	115M	493.56	602	400k	66 - 10 - 24
yelp	717K	14.0M	19.52	300	200k	75 - 10 - 15
flickr	89.2K	900K	10.09	500	70k	50 - 25 - 25

3.4 Choosing how many neighbors to sample

The variance of Poisson Sampling when $\pi_t = \frac{k}{d_s}$ is $\frac{1}{k} - \frac{1}{d_s}$. One might question why we are trying to match the variance of Neighbor Sampling and choose to use a fixed fanout for all the seed vertices. In the uniform probability case, if we have already sampled some set of edges for all vertices in S, and want to sample one more edge, the question becomes which vertex in S should we sample the new edge for? Our answer to this question is the vertex s, whose variance would improve the most. If currently vertex s has $\tilde{d}_s$ edges sampled, then sampling one more edge for it would improve its variance from $\frac{1}{\tilde{d}_s} - \frac{1}{d_s}$ to $\frac{1}{1+\tilde{d}_s} - \frac{1}{d_s}$. Since the derivative of the variance with respect to $\tilde{d}_s$ is monotonic, we are allowed to reason about the marginal improvements by comparing their derivatives, which is:

$$\frac{\partial(\frac{1}{\tilde{d}_s} - \frac{1}{d_s})}{\partial \tilde{d}_s} = -\frac{1}{\tilde{d}_s^2} \tag{19}$$

Notice that the derivative does not depend on the degree d_s of the vertex s at all, and the greater the magnitude of the derivative, the more improvement the variance of a vertex gets by sampling one more edge. Thus, choosing any vertex s with least number of edges sampled would work for us, that is: $s = \arg\min_{s' \in S} \tilde{d}_{s'}$. In the light of this observation, one can see that it is optimal to sample an equal number of edges for each vertex in S. This is one of the reasons LADIES is not efficient with respect to the number of edges it samples. On graphs with skewed degree distributions, it samples thousands of edges for some seed vertices, which contribute very small amounts to the variance of the estimator since it is already very close to 0.

4 Experiments

In this section, we empirically evaluate the performance of each method in the node-prediction setting on the following datasets: reddit (Hamilton et al., 2017), products (Hu et al., 2020), yelp, flickr (Zeng et al., 2020). More details about these datasets are given in Table 1. We compare LABOR variants LABOR-0, LABOR-1 and LABOR-*, where $0, 1, *$ stand for the number of fixed point iterations applied to optimize 13 respectively, against NS (Neighbor Sampling), LADIES and PLADIES sampling methods, where PLADIES is the unbiased Poisson Sampling variant of LADIES introduced in Section 3.1. We do not include Fast-GCN in our comparisons as it is superseeded by the LADIES paper. The works of Liu et al. (2020); Zhang et al. (2021); Huang et al. (2018); Cong et al. (2021); Dong et al. (2021) are not included in the comparisons because they either take into account additional information such as historical embeddings or their magnitudes or they have a different sampling structure such a vertex cache to sample from. Also the techniques in these papers are mostly orthogonal to the sampling problem and algorithms discussed in this paper. We evaluate all the baselines on the GCN model in Eq. 2 with 3 layers, with 256 hidden dimension and residual skip connections enabled. We use the Adam optimizer (Kingma & Ba, 2014) with 0.001 learning rate. We carried out our experiments using the DGL framework (Wang et al., 2019) with the Pytorch backend (Paszke et al., 2019)[1]. Experiments were repeated 100 times and averages are presented.

We will first show that despite the different number of sampled vertices, LABOR and NS training loss curves are almost the same in Section 4.1 with the same fanout and batch size hyperparam-

[1]Codes are in final stages of an internal review, all codes will be made publicly available to reproduce

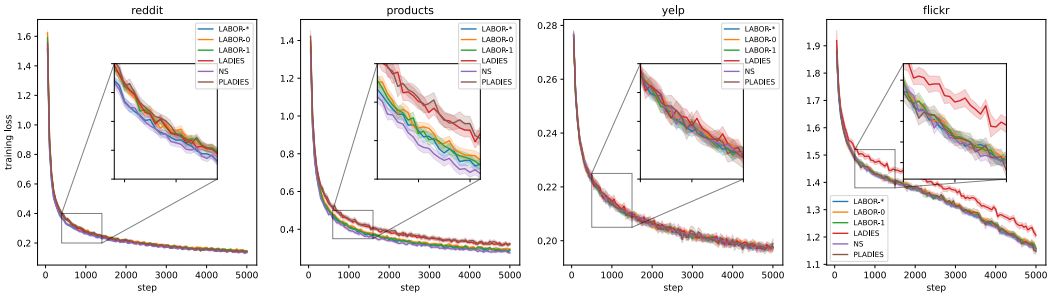

Figure 1: The loss curve on different datasets with same batch size. The soft edges represent the confidence interval. Number of sampled vertices and edges can be found in Table 2.

Table 2: Average number of vertices and edges sampled in different layers (All the numbers are in thousands, lower is better). Last column shows iterations(minibatches) per second (it/s) (higher is better). The hyperparameters of LADIES and PLADIES were picked to roughly match the number of vertices sampled by the LABOR-* to get a fair comparison. The convergence curves can be found in Figure 1. The timing information was measured on an NVIDIA T4 GPU. Green stands for best, red stands for worst results, with a 5% cutoff.

Dataset	Algo.	$\lvert V_3 \rvert$	$\lvert E_3 \rvert$	$\lvert V_2 \rvert$	$\lvert E_2 \rvert$	$\lvert V_1 \rvert$	$\lvert E_1 \rvert$	$\lvert V_0 \rvert$	it/s
reddit	PLADIES	**24.0**	**3450**	14.1	1010	5.97	35.2	1	**1.7**
	LADIES	**25.3**	**3520**	14.5	1010	5.97	34.5	1	**1.8**
	LABOR-*	**24.1**	1070	13.7	435	6.03	26.9	1	4.1
	LABOR-1	26.6	261	14.4	116	6.12	16.7	1	24.8
	LABOR-0	35.8	**177**	17.8	67.1	6.77	9.64	1	**37.6**
	NS	**167**	682	68.3	100	10.1	9.65	1	14.2
products	PLADIES	**160**	**2620**	51.2	320	9.68	12.0	1	**4.1**
	LADIES	**165**	**2560**	51.8	302	9.68	11.8	1	**4.2**
	LABOR-*	**166**	1250	51.8	167	9.78	10.6	1	6.2
	LABOR-1	178	799	53.4	136	9.78	10.5	1	21.3
	LABOR-0	237	**615**	62.4	99.6	10.1	9.89	1	**32.5**
	NS	**513**	944	95.4	106	10.6	9.89	1	24.6
yelp	PLADIES	**100**	**1370**	29.5	187	6.16	7.07	1	**5.1**
	LADIES	**102**	**1390**	29.7	191	6.17	7.08	1	**5.3**
	LABOR-*	**105**	991	30.7	158	6.15	6.83	1	13.3
	LABOR-1	109	447	31.0	96.1	6.18	6.76	1	**27.3**
	LABOR-0	138	**318**	35.1	53.9	6.25	6.29	1	**27.2**
	NS	**188**	392	42.5	54.8	6.32	6.27	1	23.0
flickr	PLADIES	**55.2**	**321**	24.9	86.0	6.23	6.93	1	**10.2**
	LADIES	**55.9**	**323**	25.1	86.2	6.23	6.93	1	**10.5**
	LABOR-*	**56.6**	308	25.6	84.7	6.29	6.95	1	20.3
	LABOR-1	57.7	242	25.9	73.4	6.29	6.93	1	**32.7**
	LABOR-0	65.9	**219**	29.1	51.9	6.37	6.71	1	**33.3**
	NS	**73.3**	244	32.8	51.9	6.37	6.72	1	**31.7**

eters. We will match the hyperparameters of LADIES with the number of vertices sampled on LABOR and see whose batches have better quality. Then, we will show what happens when different sampling algorithms are given the same budget and compare their vertex sampling efficiency in Section 4.2. Section 4.3 shows the reduction in the number of vertices sampled with each fixed point iteration. PLADIES is the poisson sampling version of LADIES the ladies algorithm as described in Section 3.1.

4.1 COMPARISON AGAINST NEIGHBOR SAMPLING AND LADIES

In this experiment, we set the batch size to $1,000$ and the fanout $k = 10$ for LABOR and NS methods to see the difference in the sizes of the sampled subgraphs and also whether convergence behaviour is the same. In Figure 1, we can see that the convergence curves of both NS and LABOR variants are almost the same showing that sampling smaller subgraphs does not really affect the batch quality. Table 2 shows the difference of the sampled subgraph sizes in each layer. One can see that on reddit, LABOR-* samples $6.9\times$ fewer vertices in the 3rd layer while keeping the same convergence behaviour. On the flickr dataset however, LABOR-* samples only $1.3\times$ fewer vertices. The amount of difference depends on two factors. The first is the amount of overlap of neighbors among the vertices in S. If the neighbors of vertices in S did not overlap at all, then one obviously can not do better than NS. The second is the average degree of the graph. With a fanout of 10, both Neighbor Sampling and LABOR has to copy the whole neighborhood of a vertex s with degree $d_s \leq 10$. Thus for such graphs, it is expected that there is a small difference because for a lot of the vertices, their whole neighborhood is copied. If we look at Table 1, the average degree of the flickr graph is 10.09, and thus there is only a small difference between LABOR and NS.

In Table 2, the number of sampled edges is another important metric to look at. We can see that LABOR-0 reduces both the number of vertices and edges sampled. On the other hand, when importance sampling is enabled, the number of vertices sampled goes down while number of edges sampled goes up. This is because when importance sampling is used, inclusion probabilities become nonuniform and it takes more edges per seed vertex to get a good approximation (see Eq. 10).

The hyperparameters of LADIES and PLADIES were picked to match LABOR-* so that all methods have the same sampling budget in each layer (see Table 2). Figure 1 shows that, in terms of loss curve, LADIES and PLADIES perform almost the same on all but the flickr dataset, in which case there is a big difference between the two in favor of PLADIES. We also see that LABOR variants either match the quality of PLADIES or outperform it, as can be seen on the products dataset.

Looking at Table 2, we can see that LABOR-0 has the best runtime performance across all datasets. This is both due to lack of the overhead of performing the fixed point iterations and also it sampled the fewest edges, compared to the other LABOR variants. By design, all LABOR variants should have the same convergence curves, as seen in Figure 1. Then, the decision of which variant to use depends on one factor: feature access speed. If vertex features were stored on a slow storage medium (such as, on host memory accessed over PCI-E), then minimizing number of sampled vertices would become the highest priority, in which case, one should pick LABOR-*. Depending on the relative vertex feature access performance and the performance of the training processor, one can choose to use LABOR-j, the faster feature access, the lower the j.

4.2 EVALUATION OF VERTEX SAMPLING EFFICIENCY

In this experiment, we set a limit on the number of sampled vertices and after each epoch, we adjust the batch size to match the given vertex budget starting with an initial batch size of 1,000. The budgets used were picked around the same magnitude with numbers in the Table 2 in the $|V_3|$ column and can be found in Table 1. Figure 2 displays the result of this experiment. We can see that the more vertex efficient the sampling method is, the larger batch size it can use during training. Number of vertices sampled is not a function of the batch size for the LADIES algorithm so we do not include it in this comparison. All of the experiments were repeated 100 times and their averages were plotted, that is why our convergence plots are smooth and differences are clear. The most striking result in this experiment is that there is a two order magnitude difference in batch sizes of LABOR-* and NS algorithms on the reddit dataset.

4.3 IMPORTANCE SAMPLING NUMBER OF FIXED POINT ITERATIONS

In this section, we look at the convergence behaviour of the fixed point iterations described in Section 3.3. Table 3 shows the number of sampled vertices in the last layer with respect to the number of fixed point iterations applied. In this table, the ∞ stands for applying the fixed point iterations until convergence, and convergence occurs in at most 15 iterations in practice before the relative change in the objective function is less than 10^{-4}. One can see that most of the reduction in the objective function 13 occurs after the first iteration, and the remaining iterations have diminishing

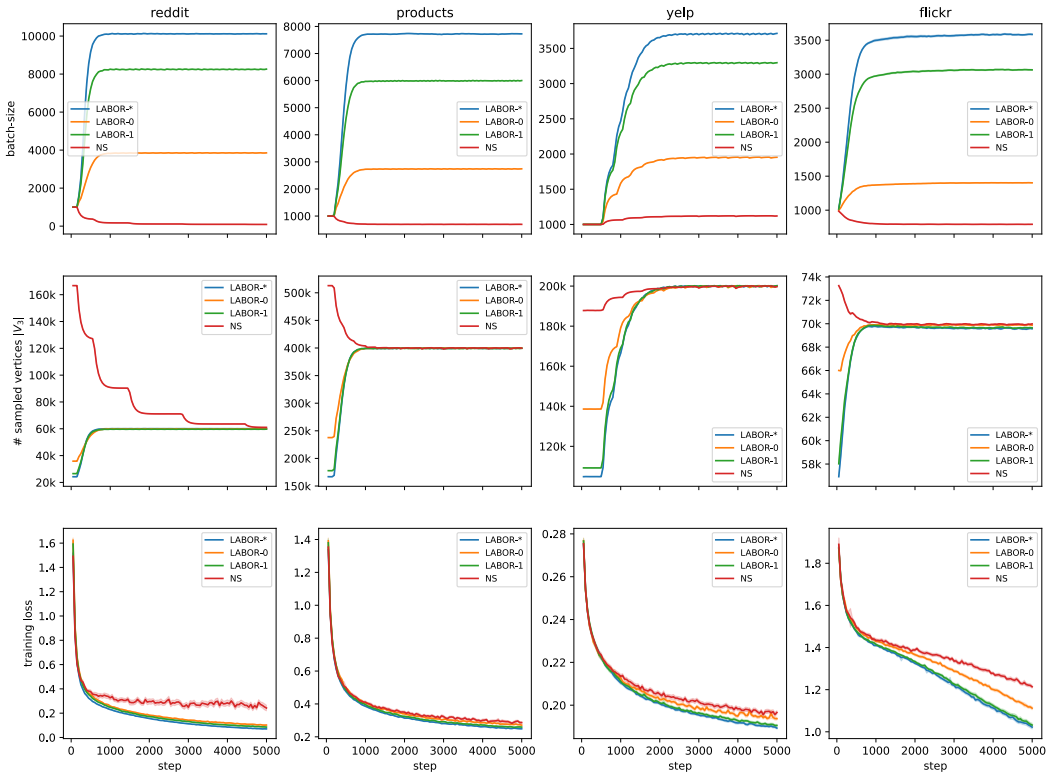

Figure 2: Vertex sampling efficiency under the same sampling budget. A starting batch size of 1k is used and the batch size is adjusted at the end of each epoch to better match the vertex budget. The first row of plots is for the batch size and the second row of plots shows the number of vertices sampled in the last layer, $|V_3|$ of running averages. The last row shows the training loss curves.

returns. Applying the fixed point iterations until convergence can save from 14% - 33% depending on the dataset. The monotonically decreasing numbers provide empirical evidence for proof in the Appendix A.1.

Table 3: Number of vertices in 3rd layer w.r.t the number of fixed point iterations (its). ∞ stands for applying the fixed point iterations until convergence, i.e., LABOR-*, 1 its stands for LABOR-1 etc. Numbers are in thousands.

Dataset	**NS**	**0 its**	**1 its**	**2 its**	**3 its**	**∞ its**
reddit	167	35.8	26.5	25.0	24.7	24.2
products	513	237	178	170	169	166
yelp	188	138	109	106	105	105
flickr	73.3	65.9	57.6	56.8	56.6	56.5

5 Conclusions

In this paper, we introduced LABOR sampling, a novel way to combine layer and neighbor sampling approaches using a vertex-variance centric framework. We then transform the sampling problem into an optimization problem where the constraint is to match neighbor sampling variance for each vertex while sampling the fewest number of vertices. We show how to minimize this new objective function via fixed-point iterations. On dense datasets like Reddit, we show that our approach can sample a subgraph with $7\times$ fewer vertices without degrading the batch quality. We also show that compared to LADIES, LABOR converges faster with same sampling budget.

REFERENCES

D.P. Bertsekas. Incremental least squares methods and the extended kalman filter. In *Proceedings of 1994 33rd IEEE Conference on Decision and Control*, volume 2, pp. 1211–1214 vol.2, 1994. doi: 10.1109/CDC.1994.411166.

Jianfei Chen, Jun Zhu, and Le Song. Stochastic training of graph convolutional networks with variance reduction. *35th International Conference on Machine Learning, ICML 2018*, 3:1503–1532, 2018.

Yifan Chen, Tianning Xu, Dilek Hakkani-Tur, Di Jin, Yun Yang, and Ruoqing Zhu. Calibrate and Debias Layer-wise Sampling for Graph Convolutional Networks. jun 2022. URL http://arxiv.org/abs/2206.00583.

Wei Lin Chiang, Yang Li, Xuanqing Liu, Samy Bengio, Si Si, and Cho Jui Hsieh. Cluster-GCN: An efficient algorithm for training deep and large graph convolutional networks. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 257–266. Association for Computing Machinery, jul 2019. doi: 10.1145/3292500.3330925.

Avery Ching, Sergey Edunov, Maja Kabiljo, Dionysios Logothetis, and Sambavi Muthukrishnan. One trillion edges: Graph processing at facebook-scale. *Proc. VLDB Endow.*, 8(12): 1804–1815, aug 2015. doi: 10.14778/2824032.2824077. URL https://doi.org/10.14778/2824032.2824077.

Weilin Cong, Morteza Ramezani, and Mehrdad Mahdavi. On the Importance of Sampling in Training GCNs: Tighter Analysis and Variance Reduction. 2021. URL http://arxiv.org/abs/2103.02696.

Jialin Dong, Da Zheng, Lin F. Yang, and George Karypis. Global Neighbor Sampling for Mixed CPU-GPU Training on Giant Graphs. *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 289–299, 2021. doi: 10.1145/3447548.3467437.

Alan H Dorfman. The Hajek Estimator Revisited. (4):760–765, 1997. URL http://www.asasrms.org/Proceedings/papers/1997_130.pdf.

Matthias Fey and Jan Eric Lenssen. Fast graph representation learning with pytorch geometric, 2019. URL https://arxiv.org/abs/1903.02428.

Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper/2017/file/5dd9db5e033da9c6fb5ba83c7a7ebea9-Paper.pdf.

C. A. R. Hoare. Algorithm 65: Find. *Commun. ACM*, 4(7):321–322, jul 1961. doi: 10.1145/366622.366647. URL https://doi.org/10.1145/366622.366647.

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. *Advances in Neural Information Processing Systems*, 2020-Decem(NeurIPS):1–34, 2020.

Wenbing Huang, Tong Zhang, Yu Rong, and Junzhou Huang. Adaptive sampling towards fast graph representation learning. *Advances in Neural Information Processing Systems*, 2018-Decem(Nips): 4558–4567, 2018.

Samir Khan and Johan Ugander. Adaptive normalization for IPW estimation. pp. 1–31, 2021. URL http://arxiv.org/abs/2106.07695.

Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2014. URL https://arxiv.org/abs/1412.6980.

Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *5th International Conference on Learning Representations, ICLR 2017 - Conference Track Proceedings*, pp. 1–14, 2017.

Ziqi Liu, Zhengwei Wu, Zhiqiang Zhang, Jun Zhou, Shuang Yang, Le Song, and Yuan Qi. Bandit samplers for training graph neural networks. *Advances in Neural Information Processing Systems*, 2020-Decem, 2020.

Xichuan Niu, Bofang Li, Chenliang Li, Rong Xiao, Haochuan Sun, Hongbo Deng, and Zhenzhong Chen. A dual heterogeneous graph attention network to improve long-tail performance for shop search in e-commerce. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '20, pp. 3405–3415, 2020. doi: 10.1145/3394486.3403393.

E Ohlsson. Sequential poisson sampling. *Journal of Official Statistics-Stockholm-*, 14(2):149–162, 1998.

Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *NeurIPS*, 2019.

Minjie Wang, Da Zheng, Zihao Ye, Quan Gan, Mufei Li, Xiang Song, Jinjing Zhou, Chao Ma, Lingfan Yu, Yu Gai, Tianjun Xiao, Tong He, George Karypis, Jinyang Li, and Zheng Zhang. Deep graph library: A graph-centric, highly-performant package for graph neural networks, 2019. URL https://arxiv.org/abs/1909.01315.

Michael S Williams, Hans T Schreuder, and Gerardo H Terrazas. Poisson Sampling – The Adjusted and Unadjusted Estimator Revisited. pp. 12, 1998.

Rex Ying, Ruining He, Kaifeng Chen, Pong Eksombatchai, William L. Hamilton, and Jure Leskovec. Graph convolutional neural networks for web-scale recommender systems. *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 974–983, 2018. doi: 10.1145/3219819.3219890.

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.

Qingru Zhang, David Wipf, Quan Gan, and Le Song. A Biased Graph Neural Network Sampler with Near-Optimal Regret. (NeurIPS):1–25, 2021. URL http://arxiv.org/abs/2103.01089.

Che Zheng, Hongzhi Chen, Yuxuan Cheng, Zhezheng Song, Yifan Wu, Changji, Li, James Cheng, Han Yang, and Shuai Zhang. Bytegnn: Efficient graph neural network training at large scale. *Proc. VLDB Endow.*, 15:1228–1242, 2022a.

Da Zheng, Xiang Song, Chengru Yang, Dominique LaSalle, and George Karypis. Distributed hybrid cpu and gpu training for graph neural networks on billion-scale heterogeneous graphs. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 4582–4591, 2022b.

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. *Advances in Neural Information Processing Systems*, 32(NeurIPS), 2019.

A Appendix

A.1 Fixed point iterations

Given any $\pi^{(0)} > 0$ and one iteration to get $\pi^{(1)}$ and one more iteration using $c(\pi^{(1)})$ to get $\pi^{(2)}$, then we have the following observation:

$$\pi_t^{(1)} = \pi_t^{(0)} \max_{t \to s} c_s(\pi^{(0)})$$
$$\frac{1}{d_s^2} \sum_{t \to s} \frac{1}{\min(1, c_s(\pi^{(0)})\pi_t^{(0)})} - \frac{1}{d_s} = \frac{1}{k} - \frac{1}{d_s}$$
$$\frac{1}{d_s^2} \sum_{t \to s} \frac{1}{\min(1, c_s(\pi^{(1)})\pi_t^{(1)})} = \frac{1}{d_s^2} \sum_{t \to s} \frac{1}{\min(1, c_s(\pi^{(1)}) \max_{t \to s'} c_{s'}(\pi^{(0)})\pi_t^{(0)})}$$

Now, note that for a given $t \in N(s)$, $\max_{t \to s'} c_{s'}(\pi^{(0)}) \geq c_s(\pi^{(0)})$ since $s \in \{s' \mid t \to s'\}$. This implies that $\max_{t \to s'} c_{s'}(\pi^{(0)})\pi_t^{(0)} \geq c_s(\pi^{(0)})\pi_t^{(0)}$. Note that for $\pi'_t = c_s(\pi^{(0)})\pi_t^{(0)}, \forall t \to s$, we have $c_s(\pi') = 1$ for a given s. Since $\max_{t \to s'} c_{s'}(\pi^{(0)})\pi_t^{(0)} \geq c_s(\pi^{(0)})\pi_t^{(0)} = \pi'_t, \forall t \to s$, this let's us conclude that the $c_s(\pi^{(1)}) \leq 1$, because the expression is monotonically increasing with respect to any of the π_t. By induction, this means that $c_s^{(i)} \leq 1, \forall i \geq 1$. Since $\pi_t^{(i)} = \pi_t^{(0)} \prod_{j=0}^{i-1} \max_{t \to s'} c_{s'}(\pi^{(i-1)})$, $\pi_t^{(i)}$ is monotonically decreasing. This means that the objective value in Eq. 13 is also monotonically decreasing and is clearly bounded from below by 0. Any monotonically decreasing sequence bounded from below has to converge, so our fixed point iteration procedure is convergent as well.

A.2 Exactly sampling a fixed number of Neighbors

One can easily resort to Sequential Poisson Sampling by Ohlsson (1998) if one wants $\tilde{d}_s = \min(k, d_s)$ instead of $E[\tilde{d}_s] = \min(k, d_s)$ to get the exact same behaviour to Neighbor Sampling. Given π_t, c_s and r_t, we pick the $\tilde{d}_s = \min(k, d_s)$ smallest vertices $t \to s$ with respect to $\frac{r_t}{c_s \pi_t}$, which can be computed in expected linear time by using the quickselect algorithm (Hoare, 1961).

A.3 Extension to the weighted case

If the given adjacency matrix A has nonuniform weights, then we want the estimate the following:

$$H_s = \frac{1}{A_{*s}} \sum_{t \to s} A_{ts} M_t \tag{20}$$

where $A_{*s} = \sum_{t \to s} A_{ts}$. The variance is as follows:

$$\mathrm{Var}(H''_s) \leq \mathrm{Var}(nH'_s) = \frac{1}{(A_{*s})^2} \Big(\sum_{t \to s} \frac{A_{ts}^2}{\min(1, c_s \pi_t)} - \sum_{t \to s} A_{ts}^2 \Big) \tag{21}$$

If we apply importance sampling from the perspective of a single vertex, then we need to make sure that the probability of picking vertex t is proportional to A_{ts}. To enforce that, we define $c_{ts} = c_s A_{ts}$. Then the variance becomes:

$$\mathrm{Var}(H''_s) \leq \mathrm{Var}(nH'_s) = \frac{1}{(A_{*s})^2} \Big(\sum_{t \to s} \frac{A_{ts}^2}{\min(1, c_{ts} \pi_t)} - \sum_{t \to s} A_{ts}^2 \Big) \tag{22}$$

In this case, we can still aim to reach the same variance target $\frac{1}{k} - \frac{1}{d_s}$ by finding c_s that satisfies the following equality:

$$\frac{1}{(A_{*s})^2} \Big(\sum_{t \to s} \frac{A_{ts}^2}{\min(1, c_{ts} \pi_t)} - \sum_{t \to s} A_{ts}^2 \Big) = \frac{1}{k} - \frac{1}{d_s} \tag{23}$$

Rearranging this expression, we get:

$$\sum_{t \to s} \frac{A_{ts}^2}{\min(1, c_{ts}\pi_t)} = (A_{*s})^2 \Big(\frac{1}{k} - \frac{1}{d_s} \Big) + \sum_{t \to s} A_{ts}^2 \tag{24}$$

One can then use the described fixed point iterations to compute both $c_s(\pi)$ for a given π, and π^* by replacing c_s with c_{ts} in 18.