Graph Learning with Loss-Guided Training

Eliav Buchnik*
eliavbuh@gmail.com
Tel Aviv University
Google Research

Edith Cohen*
edith@cohenwang.com
Google Research
Tel Aviv University

ABSTRACT

Classically, ML models trained with stochastic gradient descent (SGD) are designed to minimize the average loss per example and use a distribution of training examples that remains *static* in the course of training. Research in recent years demonstrated, empirically and theoretically, that significant acceleration is possible by methods that dynamically adjust the training distribution in the course of training so that training is more focused on examples with higher loss. We explore *loss-guided training* in a new domain of node embedding methods pioneered by DEEPWALK. These methods work with implicit and large set of positive training examples that are generated using random walks on the input graph and therefore are not amenable for typical example selection methods. We propose computationally efficient methods that allow for loss-guided training in this framework. Our empirical evaluation on a rich collection of datasets shows significant acceleration over the baseline static methods, both in terms of total training performed and overall computation.

CCS CONCEPTS

• **Computing methodologies** → **Dimensionality reduction and manifold learning.**

ACM Reference Format:

Eliav Buchnik and Edith Cohen. 2020. Graph Learning with Loss-Guided Training. In *3rd Joint International Workshop on Graph Data Management Experiences Systems (GRADES) and Network Data Analytics (NDA) (GRADES-NDA'20)*, June 14, 2020, Portland, OR, USA. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3398682.3400060>

1 INTRODUCTION

Graph data is prevalent and models entities (nodes) and interactions between them (edges). The edges may corresponds

to provided interactions between entities (likes, purchases, messages, hyperlinks) or are derived from metric data, for example, by connecting each point to its nearest neighbors.

Node embeddings, which are representations of graph nodes in the form of low dimensional vectors, are an important component in graph analysis pipelines. They are used as task-agnostic representation with downstream tasks that include node classification, node clustering for community detection, and link prediction for recommendations [5, 15, 26]. Embeddings are computed with the qualitative objective of preserving structure – so that nodes that are more connected get assigned closer embedding vectors [4, 12, 19, 22, 26, 40]. The optimization objective has the general form of a weighted sum over examples (pairs of nodes) of a per-example loss function and are commonly performed using stochastic gradient descent (SGD) [14, 19, 22, 28, 31].

1.1 Node embeddings via random walks

A first attempt to obtain positive training examples (node pairs) from the input graph is to use the provided set of edges [20]. A highly effective approach, pioneered by DEEPWALK [26], is to instead select examples based on co-occurrence of pairs in short random walks performed on the input graph. These methods weight and greatly expand the set of positive examples. DEEPWALK treats random walks on the graph as "sentences" of nodes and applies the popular word embedding framework WORD2VEC [22]. NODE2VEC [15] further refined the method by extending the family of random walks with hyper-parameters that tune the depth and breadth of the walk. Prolific followup work (see summary in [5, 37]) further extended the family of random walks but retained the general structure of producing "node sentences."

1.2 Loss-guided training

Random-walk base methods were studied in settings where the distribution of random walks and thus the distribution of training examples remain *static* in the course of training. A prolific research thread proposed methods that accelerate the training or improve accuracy by *dynamically* modifying the distribution of examples in the course of training [1, 3, 21, 32, 35, 36, 41]: These approaches include Curriculum/self-paced learning [3], where the selection is altered to mimic human learning: First the algorithm learns over the "easy"

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

GRADES-NDA'20, June 14, 2020, Portland, OR, USA

© 2020 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8021-8/20/06.

<https://doi.org/10.1145/3398682.3400060>

examples and then moves to "hard" examples, where margin is used as a measure of difficulty. A related approach guides the example selection process by the current magnitude of the gradient or the loss value. One proposed method applies importance sampling according to loss or gradient [1, 41], which preserves the expected value of the stochastic gradient updates but spreads them differently. Other methods focus on higher loss examples in a biased fashion that essentially alters the objective: Hard examples for image training [32, 36], selecting examples by moving average of the loss [21], or focusing entirely on the highest loss examples [35] with a compelling theoretical underpinning. Overall, these methods were studied with supervised learning and as far as we know, were not explored for computing node embeddings.

1.3 Our contribution

We propose and study methods that incorporate dynamic training, in particular example selection that is focused on higher loss examples, in the particular context of popular random-walk based example selection methods for node embedding. The hope is that we can obtain similar gains in performance as observed in other domains.

The application of loss-guided training to random-walk based methods poses some methodical and computational challenges. First, the methods used for other domains are not directly applicable. They were considered in supervised situations where the input data has the form of (example, label) pairs which are available explicitly and make the loss computation straightforward. In our setting, examples are produced during training using random walks: The potential number of examples can be quadratic in the number of nodes even when the input graph is sparse and the set is implicit in the graph representation. Thus, per-example state or loss evaluation on all potential examples cannot be efficiently maintained, which rules out approaches such as [21, 35].

Second, dynamic example selection, and in particular loss-guided example selection, tends to be computation heavy and trades-off the efficiency of em training (performing the gradient updates) and efficiency of *preprocessing* (the computation needed to generate the training sequence [1, 21, 41]). Even with the baseline random walk methods, the computational resources needed increase with graph size, the length and type of the random walk, the number of generated examples from the walk, and the dimension of the embedding vectors. In practice, the cost of embedding computation tends to be a significant part of the overall downstream pipeline. We aim to enhance random-walk based methods without compromising their scalability.

The components of training and preprocessing costs typically draw on different resources (e.g., gradient updates are communicated). We aim for efficiency and design loss-guided

training methods that provide tunable trade-offs. Our most effective approaches work with the same random walk processes as the respective baseline methods and assign *loss scores* to walks (each generating a set of examples) instead of to individual examples. At each selection phase we generate a set of random walks according to the baseline model, assign loss scores to these walks (via methods to be detailed later on), and choose a sample of the walks for training that is weighted by their loss scores. We empirically show that across a variety of datasets, our loss-guided methods provide dramatic reduction in training cost with a very small increase in preprocessing cost compared with the baseline methods that use a static distribution of training examples.

1.4 Related work

Graph Neural Networks (GNNs) are an emerging approach for graph learning tasks (see survey [38]). Notably, Graph Convolutional Networks [2, 11, 17] work with node features and create representations in terms of node features [40]. Variational auto-encoders [18] produce node embeddings in an unsupervised fashion but perform similarly to prior methods. Random-walk based methods remain a viable alternative that obtains state of the art results for node representations computed from graph structure alone.

1.5 Overview

The paper is organized as follows. In Section 2 we provide necessary background on the baseline node embedding methods DEEPWALK [26] and NODE2VEC [15] and the Word2Vec SGNS framework [22] that they build on. In Section 3 we present our methods that dynamically modify the distribution of training examples according to loss. We provide details on our experimental setup in Section 4. We illustrate the benefits of loss-guided training using a synthetic example network in Appendix A. The real-life datasets and tasks used in our experiments are described in Section 5 and results are reported and discussed in Section 6 and Appendix B.

2 PRELIMINARIES

We consider graph datasets of the form (V, E, w) with a set of nodes V , that represent entities, a set of edges $E \subset V \times V$ that represent pairwise interactions, and an assignment w of positive scalar weights to edges that correspond to the strength of interactions. A node embeddings is a mapping of nodes $i \in V$ to vectors $f_i \in \mathbb{R}^d$, where typically $d \ll |V|$.

2.1 Overview of baseline methods

The node embeddings methods we consider here are based on the popular DEEPWALK [26] and its refinement NODE2VEC [15]. Algorithm 1 provides a high-level view of the baseline methods. These methods build on the WORD2VEC [22] Skip

Gram with Negative Sampling (SGNS) method. SGNS was originally designed for learning embeddings for words in a text corpus. The method generates short sequences (referred to as *sentences* of consecutive words from the text corpus and uses these sentences for training (more details are provided below). The node embeddings methods generate instead sequences of nodes using short random walks on the graph and apply the SGNS framework to these node "sentences" in a black box fashion. The node embedding methods differ in the distribution over node sentences. Both our baselines specify distributions $\text{Walk}[i, t]$ of random walks of length t that start from a node i . DEEPWALK conducts a simple random walk, where the next node is selected independently of history according to weight of outgoing edges, that is, if the walk is node i then the probability of continuing to node j is $w_{ij}/\sum_h w_{ih}$. NODE2VEC uses two hyperparameters (p, q) to control the "breadth" and "depth" of the walk, in particular, to what extent it remains in the neighbourhood of the origin node. The method initializes randomly the embedding vectors and updates them according to sentences. Sentences for training are generated by selecting a start node uniformly. With both baseline methods, the distribution over sentences is *static*, that is remains the same in the course of training. To streamline the presentation, we will use the baselines as black boxes that take an input graph $G = (V, E, w)$ and for a node $i \in V$ and length $t > 0$ provide samples from the baseline-specific distribution $\text{Walk}[i, t]$.

2.2 Overview of SGNS

For completeness, we provide more details on SGNS [22]. SGNS trains two vectors for each entity i , a focus vector, f_i , and a context vector c_i .

SGNS takes as hyper parameters a "skip window" Δ and ratio λ of positive to negative examples. It works with input sentence $S := (v_0, \dots, v_t)$ as input. A sentence S is processed by generating a randomized set of pairs that are then used as positive training examples:

- (i) Draw i.i.d. $\Delta_i \sim U[t]$ for $i \in [t]$
- (ii) $\text{Pairs}(S) \leftarrow \bigcup_{i=0}^t \{(v_i, v_j) \mid |i - j| \leq \Delta_i\}$.

Skip lengths Δ_i for each v_i are selected independently uniformly at random from $\{1, \dots, t\}$. $\text{Pairs}(S)$ then includes all ordered pairs where v_j within that skip length from v_i .

For each positive example, λ random negative examples are drawn with the same focus i and a randomly selected context according to entity frequencies in positive examples to the power of 0.75. Intuitively, negative examples [16] provide an "anti-gravity" effect that prevents all embeddings from collapsing into the same vector. We denote by κ_{ij} the probability that positive example pair (i, j) is generated and

by $n_{ij} \propto \|\kappa_i\|_1 \|\kappa_j\|_1^{0.75}$ the probability that a negative example pair (i, j) is generated. The hyper parameter λ specifies a ratio of negative to positive examples. The optimization objective when using this distribution over examples has the general form:

$$L := \sum_{i,j} \kappa_{ij} L_+(i, j) + \lambda \sum_{i,j} n_{ij} L_-(i, j) . \quad (1)$$

The per-example loss functions are defined as

$$L_+(i, j) := \log(\sigma(f_i \cdot c_j)) \quad (2)$$

$$L_-(i, j) := \log(\sigma(-f_i \cdot c_j)) , \quad (3)$$

where the Sigmoid function is defined to be $\sigma(x) := \frac{1}{1+\exp(x)}$.

At a high level, the gradient updates on positive example (i, j) increase the inner product $f_i \cdot c_j$ and an update on a negative example (i, j) decreases that inner product. The SGNS objective is designed to maximize the log likelihood over all examples. This when the probability of positive example (i, j) is modeled by a Sigmoid of the inner product $\sigma(f_i \cdot c_j)$ and that of a negative example by a Sigmoid of the negated product $\sigma(-f_i \cdot c_j)$. The logarithm of the likelihood function has the form (1).

To streamline the presentation, we will treat the SGNS as a closed module for computing embedding vectors $\{f_i, c_i\}$ for $i \in V$. The module inputs V , length parameter t , and window size Δ . It has a procedure to initialize the embedding vectors. It then enters a training phase that takes as input *sentences* $S \in V^t$ and updates the embedding vectors.

3 LOSS-GUIDED TRAINING METHODS

We first discuss the challenges and design goals for loss-guided training in the SGNS-based node embedding domain. Methods in prior work were designed for supervised learning, where examples are labeled. In our setting, the SGNS loss (Equation 1) has both positive examples (that are generated from pairs co-occurring in random walks) and negative examples (that are selected randomly according to the distribution of positive examples). The negative examples distribution is therefore determined by the positive example distribution. Hence, in our setting the knob we modify would only be the distribution of positive examples.

Most methods in prior work compute (or track) approximate loss values for all examples. In our setting, the set of potential positive examples is very large, can be quadratic in the representation of the input graph dataset, and these examples are generated rather than provided explicitly. Therefore, having to maintain even approximate loss values for all potential positive examples is not feasible and can severely impact efficiency. We will instead aim to draw subsets of examples and select from these subsets according to current loss values.

Finally, the baseline methods we build on do not work with examples individually but instead generate random walks and multiple examples $\text{Pairs}(S)$ from each walk S . Using random walks rather than individual edges proved to be hugely beneficial and we do not want to lose that advantage in our loss-guided methods. Therefore, our loss-guided selection methods stick to the paradigm of generating random walks S and training with $\text{Pairs}(S)$.

3.1 Loss-guided random walks

Perhaps the most natural method to consider is to incorporate the loss values of edges in the random walks. As in the baseline methods, the start node $v_0 \sim U[|V|]$ is selected uniformly at random. A walk $(v_0, \dots, v_t)$ of length t is then computed so that

$$\Pr(v_i | v_{i-1}) = \frac{w_{v_{i-1}, v_i} L_+(v_{i-1}, v_i)^p}{\sum_{u | (v_{i-1}, u) \in E} w_{v_{i-1}, u} L_+(v_{i-1}, u)^p}, \quad (4)$$

where $p \geq 0$ is a hyper-parameter that tunes the dependence on loss. A choice of $p = 0$ provides the basic random walks used in *DeepWalk* [26]. A large value of p will result in always selecting the highest-loss outgoing edge. A value of $p = 1$ will select an edge proportionally to the product of its weight and loss value. A drawback of this method is that it is less efficient computationally: When the random walk distribution is static we can preprocess the graph so that walk generation is very efficient. Here we need to recompute loss values and edge probabilities of all outgoing edges while generating the walk. We observed empirically, however, that its performance in terms of training computation (per number of walks used for training) on almost all datasets is generally unstable and poor. This prompted us to consider instead *loss-guided selection* of walks, where the candidate random walks for training are generated as in the baseline method but the selection of walks is made according to assigned *loss scores*.

3.2 Loss-score based selection of walks

We propose a design that addresses the general issues and those specific to our settings. At a high level, we use the same random walk distribution and update and training procedures as the baseline methods (see Algorithm 1) but we modify the selection of walks for training. Algorithm 2 is a meta-algorithm for our loss-guided walk selection methods. The pseudo-code treats components as "black-boxes:" (i) The random walk distribution $\text{Walk}[i, t]$ generated from a graph $G = (V, E, w)$ according to a random process, specified start node $i \in V$ and specified length t . (ii) A training algorithm Train (such as a variant of SGNS) that includes an initialization method Train.Initialize of the embedding vectors

$\{f_i, c_i\}$ and an update method Train.Update that inputs sentences (walks) S , generates from them positive and negative examples (according to parameters on example generation Δ and λ), and performs the respective parameter updates. A component that is used only with the loss-guided methods is a *loss scoring* function $\text{Lscore}(S)$ of walks. Our choice of functions will be detailed later but they depend on specified power $p > 0$ and may also depend on the specifics of example generation from walks (see $\text{Pairs}(S)$ in Section 2.2).

For training, we initialize the embedding vectors and then repeat the following *rounds*: We draw random walks $S_i \sim \text{Walk}[i, t]$, one generated for each node $i \in V$, we compute loss scores $\text{Lscore}(S_i)$ for each of those $|V|$ walks. We then select a subset of these walks for training in a way that is biased towards the walks with the higher loss score. Specifically we will use an integer parameter $F \geq 1$ and select for training $|V|/F$ of the scored walks. The selection within each round is done using a weighted sampling without replacement method according to the loss scores $\text{Lscore}(S_i)$ of the walks $i \in V$. The weighted sampling can be implemented very efficiently in a single distributed pass over walks using each one of a variety of known order/bottom- k /varopt sampling methods (e.g. example [6–9, 13, 24, 29]). Finally, the selected walks from the round are handed to the training algorithm.

The meta procedure selects in each of $\#\text{epochs} * F$ rounds a set of $|V|/F$ walks. Therefore, selecting a total of $\#\text{epochs} * |V|$ walks in total for training. In order to compare with the baseline methods that select $|V|$ walks per epoch (one from each node), we use the term epoch to refer to providing $|V|$ walks for training. The pseudocode lists parameters that are used in various "black box" components: The length t of the generated walks, window size Δ used to generate positive examples $\text{Pairs}(S)$ from a walk S , and a power $p > 0$ which we will use later as a parameter in the scoring of walks are passed to the respective components.

Algorithm 1: Baseline method

Input: Graph $G = (V, E, w)$; Random Walk dist $\text{Walk}[i, t]$ over V^t ; walk length t ; #epochs;
 Training method Train (that uses window Δ and negatives ratio λ)

```

Train.Initialize(V)           // Initialize embedding
parameters  $\{f_i, c_i\}$  for  $i \in V$ 
foreach  $r \in [\#\text{epochs}]$  do // Select walks and update
  foreach  $v \in V$  (shuffled) do
    Draw  $S_v \sim \text{Walk}[v, t]$ 
    Train.Update( $S_v$ )         // Train on Walk  $S_v$  with
                               window size  $\Delta$ 
return Embedding vectors  $\{f_i, c_i\}$  for  $i \in V$ 

```

Algorithm 2: Walk selection by loss score

Input: Graph $G = (V, E, w)$; Random Walk dist Walk $[i, t]$ over V^t ; walk length t ; #epochs; F (fraction of walks selected per round); Walk scoring function Lscore (that uses power $p \geq 0$); Training method Train (that uses window Δ and negatives ratio λ)

Train.Initialize(V) // Initialize embedding parameters $\{f_i, c_i\}$ for $i \in V$

foreach round $\in [\text{\#epochs} * F]$ **do** // training rounds

foreach $v \in V$ **do** // Draw and score $|V|$ walks

 Draw $S_v \sim \text{Walk}[v, t]$

$L_v \leftarrow \text{Lscore}(S_v)$

$D \leftarrow$ A weighted sample without replacement of size $|V|/F$ from V according to weights L_v

foreach $v \in D$ (shuffled) **do** // Train on walks D

 Train.Update(S_v)

return Embedding vectors $\{f_i, c_i\}$ for $i \in V$

3.3 Loss scoring of walks

We consider several ways to assign loss scores to a walk $S := (v_1, \dots, v_t)$ and respective Pairs(S). All methods use a power parameter p . Our first scoring function uses the average loss of all positive examples generated from walk S :

$$\text{Lscore}_{\text{all}}(S) := \sum_{(i,j) \in \text{Pairs}(S)} L_+(i, j)^p . \quad (5)$$

The second function heuristically scores a walk by its first $t' \in [t]$ edges

$$\text{Lscore}_{t'}(S) := \sum_{i=1}^{t'-1} L_+(v_i, v_{i+1})^p . \quad (6)$$

With $t' = 1$, the walk is scored by its first edge $\text{Lscore}_1(S) := L_+(v_1, v_2)^p$.

The advantage of the loss score $\text{Lscore}_{t'}$ over $\text{Lscore}_{\text{all}}$ is that we can compute the loss scores for a candidate walk S_i from a node i without explicitly computing the walk: It suffices to draw only the first t' edges $S_{i1}, \dots, S_{it'}$ of a walk. If the node i is selected to the sample D , only then we can sample the remaining edges of the walk S_i , conditioned on its prefix being $S_{i1}, \dots, S_{it'}$. Since we compute loss scores to F times many walks than we actually train with, this is considerable saving in our per-round preprocessing cost. The disadvantage of the loss score (6) is that we are only using t' examples from the set Pairs(S) to determine the loss score, so we can expect a reduction in effectiveness.

The power p in the computation of loss scores has the role of a hyperparameter: High values of p focus the selection more on walks with high loss examples whereas lower values allow for a broader representation of walks in the training. Interestingly, Shalev-Shwartz and Wexler [35] considered the more extreme objective of minimizing the maximum

per-example loss. This objective is very sensitive to outliers (persistent high loss examples) and in some cases can divert all the training effort to be futilely spent on erroneous examples. Note that in our setting, we are not as exposed because the walks pool we select from in each round is randomized and we use without replacement sample to select a $1/F$ fraction of that pool for training.

3.4 Complexity analysis

The per-epoch training cost with both the baseline and loss-guided selection methods amounts to computing the gradients of the loss functions $L_+(i, j)$ or $L_-(i, j)$ (for positive and negative examples) and applying gradient updates. The training cost is proportional to the total number of examples generated from $|V|$ walks. The expected number of positive examples, $E[|\text{Pairs}(S)|]$, depends on the walk length t and window Δ . The total number also depends on the negatives to positives ratio λ (see Section 2.2). Therefore, each walk S generates in expectation $E[|\text{Pairs}(S)|](\lambda + 1)$ training examples. We train on $|V|$ walks in each epoch and thus the per-epoch training cost is:

$$C_{\text{train}} := E[|\text{Pairs}(S)|](\lambda + 1)|V| . \quad (7)$$

We next consider the per-epoch total computation cost, that includes preprocessing cost. For the baseline methods, the preprocessing cost corresponds to generating $|V|$ random walks (t edge traversals each). The total cost is dominated by the gradient computations of the training cost and is:

$$C_{\text{baseline}} = |V|(\lambda + 1) E[|\text{Pairs}(S)|] . \quad (8)$$

For the loss-guided methods, the preprocessing cost involves evaluations of the loss $L_+(i, j)$ on positive examples. With loss score $\text{Lscore}_{t'}$ in each round we generate the first t' steps of a random walk from each node $i \in V$. We then evaluate the loss score for each of the walks, which amounts to evaluating $L_+(i, j)$ on t' pairs (only $|V|/F$ of the walks are selected for training). The total number of loss evaluation per epoch (F rounds) is $C_{\text{prep}, t'} = F|V|t'$. With the loss score $\text{Lscore}_{\text{all}}$ we generate in each round a complete walk from each node and evaluate the $L_+(i, j)$ for each pair in Pairs(S). The total number of loss evaluations is $C_{\text{prep}, \text{all}} := E[|\text{Pairs}(S)|]F|V|$. The total computation cost combines the training and preprocessing cost and is measured by the number of loss or gradient evaluations. Note that loss or gradient evaluations have similar complexity and amount to computing $e^{\langle x, y \rangle}$ for loss and $x e^{\langle x, y \rangle}$ for the gradient. Summarizing, we have

$$C_{t'} = |V|Ft' + (\lambda + 1) E[|\text{Pairs}(S)|] = |V|Ft' + C_{\text{baseline}} \quad (9)$$

$$C_{\text{all}} = |V|(F + \lambda + 1) E[|\text{Pairs}(S)|] = |V|F E[|\text{Pairs}(S)|] + C_{\text{baseline}} . \quad (10)$$

4 EMPIRICAL EVALUATION SETUP

As baseline and for walk generation with our methods we used DEEPWALK [26] and NODE2VEC [15]. These methods define the random walk distributions $\text{Walk}[i, t]$. When evaluating our methods we fit hyperparameters to the respective baseline method. With NODE2VEC we searched over values $p, q \in \{0.25, 0.5, 1, 2\}$.

We trained models using the Gensim package [27] that builds on a classic implementation of SGNS [34]. We used the default parameters that generally perform well: $t = 10$ for the length of the walk (sentence), $\Delta = 10$ for window size, and $\lambda = 5$ for the number of negative examples generated for each positive example. With these values, we have in expectation $E[|\text{Pairs}(S)|] = 63$ positive examples generated from each walk and $E[|\text{Pairs}(S)|](\lambda + 1) = 380$ examples in total generated for each walk when training.

In our implementation we applied the baseline method (Algorithm 1) for the first epoch (training on $|V|$ walks $S_i \sim \text{Walk}[i, t]$) and applied loss-guided methods (Algorithm 2) starting from the second epoch. This is because we expect scoring by loss to not be helpful initially, with random initialization. We used $F \in \{2, 5, 10, 20\}$ rounds per epoch and power of the loss value $p \in \{1, 4, 32\}$. Each experiment is repeated r times and we report the average quality and standard error. We fit parameters on one dataset from each collection using one set of repetitions and use the same parameters with all datasets and a fresh set of repetitions.

As mentioned, SGNS determines the distribution of negative examples according to the frequencies of words in the provided positive examples. With the baseline methods, the distribution of random walks and hence the frequencies of words in positive examples remain fixed throughout training and are approximated by maintaining historic counts from the beginning of training. With our loss-guided selection, the distribution of positive examples changes over time. We experimented with different variations that use a recent positive distribution (per-round or for few recent epochs) to guide the negative selection. We did not observe significant effect on performance and report results with respect to frequencies collected since the start of training.

4.1 Tasks and metrics

We evaluated the quality of the embeddings on the following tasks, using corresponding quality metrics:

- *Clustering*: The goal is to partition the nodes into k clusters. The embedding vectors are used to compute a k -means clustering of nodes. We used `sklearn.cluster.KMeans` from `scikit-learn` package [25] with default parameters. Our quality measure is the *modularity score* [23] of the clustering.
- *Multi-class (or multi-label) classification*: Nodes have associated classes (or labels) from a set L . The class (or the set

of labels) are provided for some nodes and the goal is to learn the class/labels of remaining nodes. An embedding is computed for all nodes (in an unsupervised fashion). Following that, a supervised learning algorithm is trained on embedding and class/label pairs. We used One-vs-Rest logistic regression from the `scikit-learn` package with default parameters [25]. For multi-labels we used the multinomial option. In a multi-class setting, we obtain a class prediction from the embedding vector for each of the remaining nodes and report the fraction of correct predictions. In a multi-label setting, we provide the number of labels and the embedding vector and obtain a set of predicted labels for each node. We report the micro-averaged F1 score.

4.2 Measuring gain

Across our datasets, peak accuracy with loss-guided selection was equal or higher than baseline. We thus consider efficiency, which we measure using $\text{\#epochs}(\text{method})$, the average number of training epochs over repetitions needed for the method to reach 0.95 of peak accuracy. We can now express the training and computation cost and respective gains. With the parameter values we use, the per-epoch training cost is $C_{\text{train}} = 380|V|$ and the per-epoch computation cost is $C_{t'} = |V|(Ft' + 380)$ and $C_{\text{all}} = |V|(F + 6)63$. Accordingly, we express the gain of a loss-guided method with scoring function $\text{Lscore}_{\text{method}}$ with respect to the baseline:

- *Training gain*: is the relative decrease in number of training epochs (recall that training cost per epoch is similar for all methods).

$$1 - \frac{\text{\#epochs}(\text{method})}{\text{\#epochs}(\text{baseline})} . \quad (11)$$

When reporting the training gain, we report the error over repetitions: We compute the (sample) standard deviation of the number of epochs used by the method to reach peak (over repetitions) and normalize it by dividing by $\text{\#epochs}(\text{baseline})$.

- *Computation gain*: The relative decrease in computation cost

$$1 - \frac{C_{\text{method}} \cdot \text{\#epochs}(\text{method})}{C_{\text{train}} \cdot \text{\#epochs}(\text{baseline})} . \quad (12)$$

With $\text{Lscore}_{t'}$ the computation gain is $1 - \frac{(380+Ft') \cdot \text{\#epochs}(t')}{380 \cdot \text{\#epochs}(\text{baseline})}$ and with $\text{Lscore}_{\text{all}}$ it is $1 - \frac{(380+63F) \cdot \text{\#epochs}(\text{all})}{380 \cdot \text{\#epochs}(\text{baseline})}$.

5 DATASETS AND TASKS

We evaluate our methods on three collections of real-world datasets, summarized in Table 1. The datasets have different learning tasks (see Section 4.1):

- *Facebook page networks (clustering)*: The collection represent mutual "like" networks among verified Facebook pages. There are six datasets for different communities (TV shows,

dataset	$ V $	$ E $	
[30] Facebook pages			clustering k
Athletes	13,866	86,859	20
Company	14,113	52,311	20
Government	7,057	89,456	20
New Sites	27,917	206,259	20
Politicians	5,908	41,730	20
Public Figures	11,565	67,115	20
TV Shows	3,892	17,263	20
[33] Citation networks			Multi-class $ L $
Pubmed	19,717	44,338	3
Cora	2,708	5,429	7
Citeseer	3,327	4,732	6
[15] Protein Interactions			Multi-label $ L $
PPI	3890	38739	50

Table 1: Summary of dataset properties and tasks.

athletes, and more) [30]. The task (following [30]) is to compute embeddings with $d = 16$ and cluster the data to $k = 20$ clusters. .

- *Citation networks (multi-class)*: The collection has three networks (Citeseer, Cora and Pubmed) [33]. Networks are formed by having a node for each document and an (undirected, unweighted) edge for each citation link. Each document has a class label. Following [15, 39], we train a $d = 128$ dimensional embedding and use a random selection of 20 nodes per class as labeled training examples.
- *Protein-Protein Interactions (PPI) (multi-label)*: The dataset is a graph of human Protein-Protein interactions [15]. Each protein (node) has multiple labels and the goal is to predict this set of labels. Following [15], we use $d = 128$ and use 50% of nodes (selected uniformly at randomly) for training.

6 EMPIRICAL RESULTS

We evaluate our methods using three key metrics: Quality, training gain, and computation gain. We use figures to show quality in the course of training: We plot average performance over repetitions and provide error bars that correspond to one SD. We use tables to report training and computation gains for different methods and hyper-parameter settings. In Appendix B we provide parameter sweeps on the number of rounds per epoch F , the loss power p , and t' in the loss score $L_{score_{t'}}$. In this section we report results for $F = 10$ rounds per epoch, which seems to be a sweet spot for the training cost (see Appendix B for other values of F). We use both DEEPWALK and NODE2VEC baselines. For each loss scoring function we used the best performing overall power p : L_{score_1} performed well with $p = 32$ (selecting the $|V|/10$ highest loss walks in each round). $L_{score_{10}}$ performed well with $p = 4$ (weighted sampling that is biased towards higher loss). Interestingly, $L_{score_{all}}$ did not perform better

dataset	t', p	Training		Comp	Training		Comp
		%gain	%SD	%gain	%gain	%SD	%gain
DEEPWALK baseline				NODE2VEC baseline			
Athletes	1, 32	12	1.8	9.8	12.91	2.40	10.7
	10, 4	18.2	3.10	0.70	18.2	2.33	0.14
Company	1, 32	18.2	1.86	16.0	20.0	2.21	17.7
	10, 4	22.3	1.71	5.5	22.6	1.50	5.38
Government	1, 32	10.7	2.67	8.47	10.4	2.10	8.13
	10, 4	21.9	1.90	5.60	20.3	2.10	2.61
New Sites	1, 32	15.2	3.5	10.1	12.5	3.63	10.5
	10, 4	4.21	9.58	-17.7	7.08	8.70	-14.5
Politicians	1, 32	17.9	2.19	15.8	18.2	2.51	16.0
	10, 4	24.6	1.60	9.44	24.1	1.73	8.16
Public Figures	1, 32	10.2	5.08	7.93	7.10	3.81	4.9
	10, 4	24.3	2.97	7.69	23.3	1.84	6.05
TV Shows	1, 32	21.76	1.14	19.65	21.6	1.63	20.0
	10, 4	25.2	1.51	9.88	24.4	1.3	8.88

Table 2: Facebook Pages: Training and computation gain of loss-guided with DEEPWALK and NODE2VEC ($(p, q) = (2, 1)$) baselines ($r = 200$ repetitions, $F = 10$ rounds per epoch, ($L_{score_1}, p = 32$) and ($L_{score_{10}}, p = 4$))

than $L_{score_{10}}$ even in terms of training cost (since it is computation heavy, there is also no improvement in computation cost). We show performance with $L_{score_{all}}$ in plots but do not report it in tables.

6.1 Facebook networks (clustering task)

Representative results based on $r = 200$ repetitions are reported in Table 2 for both baselines. Figure 1 shows the modularity score in the course of training for representative datasets and methods. We fitted the Node2vec parameters to $p = 2$ and $q = 1$ on the Athletes dataset and applied with all datasets and methods. We see that loss-based selection obtained 13%-25% reduction in training and 6%-20% reduction in computation for both baselines. We can see that on almost all datasets in this collection $L_{score_{10}}$ outperformed L_{score_1} in terms of training cost but in most cases had a lower overall gain in computation cost.

6.2 Citation networks (Multi-class)

Representative results with $r = 400$ repetitions are reported in Table 3. Figure 2 shows performance in the course of training for the Pubmed dataset. The NODE2VEC parameters were fitted on the Cora dataset to $(p, q) = (2, 0.25)$. We can see that the loss-guided methods had training gains of 8%-12% on the Pubmed and Citeseer datasets, but due to large error bars there is no significance for the improvements on Cora. The loss score L_{score_1} outperformed others also in terms of training cost.

Figure 1: Modularity score in the course of training on the Athletes, Company, and Politicians networks from the Facebook collection. We show DEEPWALK and loss-guided with different walk scoring functions.

Figure 2: Pubmed dataset: Accuracy in the course of training for DEEPWALK and loss-guided methods with different scoring functions.

Figure 3: PPI dataset: Averaged micro-F1 score in the course of training for DEEPWALK and loss-guided methods with different scoring functions.

6.3 PPI network (multi-label)

Representative results with $r = 200$ repetitions with loss scores $Lscore_{t'}$ for $t' \in \{1, 3, 10\}$ are reported in Table 3 and Figure 3. NODE2VEC parameters were fitted to $(p, q) = (1, 2)$. We observe that training costs improves with t' , and in particular the training gain with $Lscore_{10}$ is significantly higher than with $Lscore_1$, but most of that gain is already attained by $Lscore_3$. The computation gain is largest with $Lscore_3$, which attains nearly the same training gain as $Lscore_{10}$ but

dataset	t', p	Training		Comp	Training		Comp
		%gain	%SD	%gain	%gain	%SD	%gain
DEEPWALK baseline							
Citation Networks, $r = 400$							
Pubmed	1, 32	9.07	3.91	7.38	9.06	3.55	8.28
	10, 4	5.21	7.02	-10.5	6.14	5.96	-10.7
Cora	1, 32	1.80	8.60	0.00	4.08	7.45	4.23
	10, 4	5.20	8.12	-12.4	8.27	6.24	-9.84
Citeseer	1, 32	7.64	6.60	5.81	11.57	5.43	9.6
	10, 4	5.73	6.20	-11.2	7.90	8.37	-9.90
Protein Interaction Network, $r = 200$							
PPI	1, 32	12.7	3.91	3.90	10.4	7.82	10.7
	3, 32	20.7	3.77	14.75	21.4	3.73	11.8
	10, 4	22.2	3.38	4.06	22.2	3.90	4.50

Table 3: Citation and PPI Networks: Training and computation gain of loss-guided selection with DEEPWALK and NODE2VEC baselines. $F = 10$, ($Lscore_1, p = 32$), ($Lscore_{10}, p = 4$) and for the PPI network also ($Lscore_3, p = 32$).

at lower per-epoch computation. Overall, we see training gains of 22% and computation gains of 12%-15%.

7 CONCLUSION

We study loss-guided example selection, known to accelerate training in some domains, for methods such as DEEPWALK and NODE2VEC that learn node embeddings using random walks. The random walk base methods use a static distribution over an implicitly-represented extended set of training examples and seems less amenable for dynamic loss-guided example selection. We propose efficient methods that facilitate loss-based dynamic example selection while retaining the highly effective structure of random walks and scalability. We demonstrate empirically the effectiveness of the proposed methods. An interesting open question is to explore such benefits with other frameworks that generate training examples on-the-fly from an implicit representation such as example augmentation or together with methods that work with feature representation of nodes such as Pinsage [40].

Acknowledgements This research is partially supported by the Israel Science Foundation (Grant No. 1595/19). We thank the anonymous GRADES-NDA '20 reviewers for many helpful comments.

REFERENCES

- [1] G. Alain, A. Lamb, C. Sankar, A. C. Courville, and Y. Bengio. Variance reduction in SGD by distributed importance sampling. *CoRR*, abs/1511.06481, 2015.
- [2] J. Atwood and D. Towsley. Diffusion-convolutional neural networks. In *NIPS*, 2016.
- [3] Y. Bengio, J. Louradour, R. Collobert, and J. Weston. Curriculum learning. In *ICML*, 2009.
- [4] Michael W. Berry, Susan T. Dumais, and Gavin W. O'Brien. Using linear algebra for intelligent information retrieval. *SIAM Review*, 37(4):573–595, 1995.
- [5] HongYun Cai, Vincent W. Zheng, and Kevin Chen-Chuan Chang. A comprehensive survey of graph embedding: Problems, techniques, and applications. *IEEE Trans. Knowl. Data Eng.*, 30(9):1616–1637, 2018.
- [6] M. T. Chao. A general purpose unequal probability sampling plan. *Biometrika*, 69(3):653–656, 1982.
- [7] E. Cohen, N. Duffield, C. Lund, M. Thorup, and H. Kaplan. Efficient stream sampling for variance-optimal estimation of subset sums. *SIAM J. Comput.*, 40(5), 2011.
- [8] E. Cohen and H. Kaplan. Summarizing data using bottom-k sketches. In *ACM PODC*, 2007.
- [9] E. Cohen and H. Kaplan. Sketch-based estimation of subpopulation-weight. Technical Report 802.3448, CORR, 2008.
- [10] Anne Condon and Richard M. Karp. Algorithms for graph partitioning on the planted partition model. *Random Struct. Algorithms*, 18(2):116–140, 2001.
- [11] M. Defferrard, X. Bresson, and P. Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *NIPS*, 2016.
- [12] J. Deng, W. Dong, R. Socher, L-J Li, K. Li, and F-F Li. Imagenet: A large-scale hierarchical image database. In *In CVPR*, 2009.
- [13] N. Duffield, M. Thorup, and C. Lund. Priority sampling for estimating arbitrary subset sums. *J. Assoc. Comput. Mach.*, 54(6), 2007.
- [14] R. Gemulla, E. Nijkamp, P. J. Haas, and Y. Sismanis. Large-scale matrix factorization with distributed stochastic gradient descent. In *ACM KDD 2011*. ACM, 2011.
- [15] A. Grover and J. Leskovec. node2vec: Scalable feature learning for networks. In *KDD*. ACM, 2016.
- [16] Y. Hu, Y. Koren, and C. Volinsky. Collaborative filtering for implicit feedback datasets. In *ICDM*, 2008.
- [17] T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. In *ICLR*, 2017.
- [18] Thomas N. Kipf and Max Welling. Variational graph auto-encoders, 2016.
- [19] Y. Koren. Factorization meets the neighborhood: a multifaceted collaborative filtering model. In *KDD*, 2008.
- [20] Y. Koren, R. Bell, and C. Volinsky. Matrix factorization techniques for recommender systems. *Computer*, 42, 2009.
- [21] I. Loshchilov and F. Hutter. Online batch selection for faster training of neural networks. In *ICLR*, 2016.
- [22] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean. Distributed representations of words and phrases and their compositionality. In *NIPS*, 2013.
- [23] M. E. J. Newman and M. Girvan. Finding and evaluating community structure in networks. *Phys. Rev. E*, 69(2), February 2004.
- [24] E. Ohlsson. Sequential poisson sampling. *J. Official Statistics*, 14(2):149–162, 1998.
- [25] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [26] B. Perozzi, R. Al-Rfou, and S. Skiena. Deepwalk: Online learning of social representations. In *KDD*. ACM, 2014.
- [27] Radim Rehůrek and Petr Sojka. Software Framework for Topic Modelling with Large Corpora. In *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks*, pages 45–50, Valletta, Malta, May 2010. ELRA. <http://is.muni.cz/publication/884893/en>.
- [28] H. Robbins and D. O. Siegmund. A convergence theorem for non negative almost supermartingales and some applications. In J. S. Rustagi, editor, *Optimizing Methods in Statistics*. Academic Press, 1971.
- [29] B. Rosén. Asymptotic theory for order sampling. *J. Statistical Planning and Inference*, 62(2):135–158, 1997.
- [30] Benedek Rozemberczki, Ryan Davies, Rik Sarkar, and Charles A. Sutton. GEMSEC: graph embedding with self clustering. In *ASONAM '19: International Conference on Advances in Social Networks Analysis and Mining, Vancouver, British Columbia, Canada, 27-30 August, 2019*, pages 65–72, 2019.
- [31] R. Salakhutdinov, A. Mnih, and G. Hinton. Restricted boltzmann machines for collaborative filtering. In *ICML*. ACM, 2007.
- [32] F. Schroff, D. Kalenichenko, and J. Philbin. Facenet: A unified embedding for face recognition and clustering. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, 2015.
- [33] Prithviraj Sen, Galileo Namata, Mustafa Bilgic, Lise Getoor, Brian Gallagher, and Tina Eliassi-Rad. Collective classification in network data. *AI Magazine*, 29(3):93–106, 2008.
- [34] word2vec implementation.
- [35] S. Shalev-Shwartz and Y. Wexler. Minimizing the maximal loss: How and why. In *ICML*, 2016.
- [36] A. Shrivastava, A. Gupta, and R. B. Girshick. Training region-based object detectors with online hard example mining. In *CVPR*, 2016.
- [37] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and Philip S. Yu. A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, page 1–21, 2020.
- [38] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and Philip S. Yu. A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, page 1–21, 2020.
- [39] Z. Yang, W. W. Cohen, and R. Salakhutdinov. Revisiting semi-supervised learning with graph embeddings. In *ICML*. JMLR.org, 2016.
- [40] 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, KDD '18*. ACM, 2018.
- [41] P. Zhao and T. Zhang. Stochastic optimization with importance sampling for regularized loss minimization. In *ICML*, 2015.

A SYNTHETIC COMMUNITIES GRAPH

We start with a simple synthetic network that demonstrates the benefits of loss-guided selection. The example network structure is illustrated in Figure 4. We have three communities (red, green, and blue) of the same size. The red and green communities are inter-connected and the blue community is isolated. The goal is to reconstruct the ground truth community affiliations from the learned embedding. Our construction is inspired by random GnP graphs (each community is a GnP graph) and the planted partition model [10]. Each of the communities has 10^4 nodes. We generated intra-community edges so that each pair (i, j) of same-community

Figure 4: A community structure with three communities in green, red, and blue. The red and green communities are inter-connected and the blue community is isolated.

nodes has a connecting edge with probability $p = 0.001$. Each inter-community pair (i, j) from the red and green communities has a connecting edge with probability $q = 0.0003$.

We trained node embeddings using DEEPWALK and using loss-guided selection with DEEPWALK as a baseline. The baseline method DEEPWALK selects a start node of a random walk uniformly and hence the distribution of training examples remains balanced among the three communities through the course of training. The loss-guided selection will focus more training on walks with a higher loss score. We expect the isolated community to separate out early in training and for the two inter-connected communities to require more training to "separate" from each other. The loss $L_+(i, j)$ of a same-community pair (i, j) will be lower earlier for the isolated community. A loss-guided method after the initial stage of training is more likely to select training examples from the two inter-connected communities and thus be more productive. The benefit is further boosted by the corresponding selection of negative examples, where a community not selected for positive examples also does not participate in negative examples. The quality was measured by treating the problem as a 3-class classification problem as explained in Section 4.1 with classes assigned according to community with $L = \{\text{red, green, blue}\}$. Half the nodes (selected randomly) were used as labeled examples for the supervised component. We used $r = 25$ repetitions for each method and report representative results for embedding dimension $d = 10, 50$. Figure 5 shows the fraction of correct classifications as a function of training epochs. We observe that the different methods behave the same in the initial phase of training, until the blue community separates out from the other two but after that the loss-guided methods are more effective. The loss-score function Lscore_1 that uses the first edge of the walk attains the full advantage of the loss guided methods. This because the first edge identifies the community. Figure 6 reports the fraction of training spent at each community. We can see that in the initial phase all methods are balanced but as expected, the baseline DEEPWALK remains balanced whereas the loss-guided variant spend increasing fraction of training resources on the green and red communities, where it is more helpful.

B HYPERPARAMETER SENSITIVITY

We explore the dependence of the performance of our loss-guided methods on the following parameters: The number of edges t' out of the $t = 10$ walk edges that are used in the walk loss scoring

Figure 5: Example network: Accuracy in the course of training using DEEPWALK and loss-guided selection with DEEPWALK baseline. $d = 10$ (left) and $d = 50$ (right). The x -axis shows the number of epochs.

Figure 6: Example network: Fraction of training spent on walks starting in each community, in the course of training. DEEPWALK (left) and loss-guided (right). $d = 10$ (top) and $d = 50$ (bottom).

function $\text{Lscore}_{t'}$, the number F of rounds per epoch, and the loss power value p which determines how we weigh the loss of examples when we compute loss scores of walks.

B.1 Loss scoring methods of walks

We proposed (see Section 3.3) several loss scoring functions of walks: $\text{Lscore}_{t'}(S)$ for $t' \in [10]$ which uses the average loss of the first t' edges of the walk S and $\text{Lscore}_{\text{all}}(S)$ which uses the average loss of all positive training examples generated from the walk S . We observed empirically that $\text{Lscore}_{\text{all}}$ rarely outperformed Lscore_{10} , even in terms of training cost. We note that due to technical reasons we used the expected loss on the selected walk S (under random draws of Δ_i) instead of the precise evaluation on the pairs generated from the selected walk. This could have impacted adversely the reported performance of $\text{Lscore}_{\text{all}}$.

We explore the training and computation cost with $\text{Lscore}_{t'}$ as we vary $t' \in \{1, 3, 5, 10\}$. Representative results (with $F = 10$ and $p = 32$) are reported in Table 4 (We report results for datasets for which the error bars are small compared with the gain and its variation.) We see a general trend of improved training cost as

dataset	t'	p	F	Training		Comput
				%gain	%SD	%gain
DEEPWALK baseline						
PPI	1	32	10	12.7	3.91	10.4
	2			20.7	3.77	14.8
	5			23.3	3.90	14.0
	10			25.9	3.20	8.50
Pubmed	1	32	10	9.07	3.91	7.38
	2			7.31	5.77	4.11
	5			1.49	9.00	-6.83
	10			-10.5	19.8	-29.9
Athletes	1	32	10	12.0	1.89	9.82
	2			13.5	2.34	9.54
	5			16.8	2.60	7.66
	10			15.5	4.49	-2.62
Company	1	32	10	18.2	1.86	16.0
	2			18.4	2.06	14.5
	5			20.6	1.47	11.8
	10			20.1	3.34	2.87
Government	1	32	10	10.7	2.67	8.47
	2			12.7	3.17	8.67
	5			19.7	2.99	10.7
	10			23.6	2.53	7.11
Politicians	1	32	10	17.9	2.19	15.8
	2			19.5	2.21	15.8
	5			23.6	1.75	15.6
	10			24.8	1.38	9.72
Public figures	1	32	10	10.2	5.08	7.93
	2			12.2	4.93	8.10
	5			23.0	1.84	14.2
	10			26.9	3.41	10.9
TV shows	1	32	10	21.8	1.14	19.7
	2			22.6	1.78	18.9
	5			23.8	1.43	15.7
	10			24.4	1.55	8.82

Table 4: Varying the loss scoring function $L_{\text{score}_{t'}}$ for $t' \in \{1, 2, 5, 10\}$: Training gain and computation gain of selected datasets ($F = 10$, $p = 32$, DEEPWALK baseline).

we increase t' , but the extent of this improvement widely varies between datasets. For example, the improvement is small for the TV shows dataset, moderate for the Politicians dataset, and significant for the Government and PPI datasets. Note that the per-epoch computation cost also increases with t' (see analysis in Section 3.4 and Section 4.2). The overall computation gain as we increase t' reflects both the decrease in the number of epochs and the increase in per-epoch computation. We can see that the computation gain is often maximized for lower values of t' than the value that maximizes the training gain.

B.2 Rounds per epoch F

The parameter F controls the number of rounds per epoch. Recall that in each round we score $|V|$ walks and select $|V|/F$ of these walks for training. The setting $F = 1$ corresponds to the baseline method. In Table 5 we report training and computation gains over the DEEPWALK baseline method for $F \in \{2, 5, 10, 20\}$. We report on all datasets for which the standard deviation on the gain allowed for meaningful comparisons. Gains are reported with configurations ($L_{\text{score}_1, p = 32}$) and ($L_{\text{score}_{10}, p = 4}$). We highlight the value that maximizes the training cost or computation cost for each configuration. We can see a trend where the training gain increases with F . We see that the computation gain is often maximized at a lower F value than the value that maximizes the training gain. This

dataset	F	Training		Comp	Training		Comp
		%gain	%SD	%gain	%gain	%SD	%gain
DEEPWALK baseline							
Lscore ₁ , $p = 32$							
PPI	2	10.4	5.30	9.71	16.7	5.20	12.4
	5	13.2	5.30	11.8	20.0	4.80	10.4
	10	12.7	3.91	10.4	22.3	3.38	4.06
	20	10.1	7.50	5.70	21.5	3.40	-14.7
Athletes	2	7.04	4.30	6.46	10.4	4.80	6.35
	5	11.7	2.57	10.5	17.5	3.13	8.45
	10	12.0	1.89	9.82	18.2	3.09	0.70
	20	9.93	2.51	5.85	18.6	2.11	-15.8
Company	2	12.0	2.93	11.3	13.2	2.68	9.14
	5	19.2	2.20	17.8	20.0	2.11	11.1
	10	18.2	1.86	16.1	22.3	1.71	5.54
	20	18.0	1.69	14.1	22.2	1.26	-10.8
Government	2	6.93	3.19	6.35	15.3	4.06	11.3
	5	10.5	3.31	9.27	20.3	2.47	11.4
	10	10.7	2.67	8.47	22.0	1.89	5.06
	20	9.75	2.64	5.60	22.4	1.86	-10.7
New Sites	2	13.3	2.95	12.6	15.6	5.65	11.4
	5	17.1	3.48	15.7	8.30	7.58	-2.40
	10	15.2	3.53	12.86	4.21	9.58	-17.8
	20	11.1	4.60	6.81	4.33	10.6	-39.5
Politicians	2	11.2	4.55	10.6	14.6	2.00	10.8
	5	16.8	2.38	15.6	21.9	2.14	13.7
	10	17.9	2.19	15.8	24.6	1.60	9.44
	20	19.1	1.83	15.4	25.2	0.88	-4.08
Public Figures	2	5.63	7.39	5.07	15.5	4.58	11.5
	5	10.8	4.20	9.56	21.6	4.19	12.7
	10	10.2	5.08	7.93	24.3	2.97	7.69
	20	7.66	3.12	3.41	26.3	1.98	-5.30
TV Shows	2	10.5	3.09	9.88	17.4	3.59	13.6
	5	20.7	1.40	19.4	23.3	1.72	15.1
	10	21.8	1.14	19.7	25.2	1.51	9.90
	20	21.8	1.20	18.2	26.1	1.34	-3.40

Table 5: Varying the number of rounds per epoch $F \in \{2, 5, 10, 20\}$: Training gain and computation gain of selected datasets with respect to DEEPWALK baseline. For loss-guided with ($L_{\text{score}_1, p = 32}$) and ($L_{\text{score}_{10}, p = 4}$).

is because the per-epoch computation also increases with F (see Section 3.4 and Section 4.2) (12)). For some F value we reach a sweet spot that balances the benefits from reduced training (that increase with F) and the higher per-epoch pre-processing computation (that increases with F).

Qualitatively, higher F values mean that the training is more focused on high loss examples and that the loss values are more current. This is helpful to some point, but with high enough F we might direct all the training to outliers or erroneous examples. In the table we do not see a point where the training cost starts increasing with F but we do see that there is almost no gain between $F = 10$ and $F = 20$.

B.3 Loss power p

The loss power p is used in the walk loss scoring function (see Section 3.3). The value of p determines to what extent the training selection made in each round is biased towards examples with

dataset	p	Training		Comp	Training		Comp
		%gain	%SD	%gain	%gain	%SD	%gain
		Lscore ₁			Lscore ₁₀		
PPI	1	6.02	4.33	3.71	10.9	3.80	-9.70
	4	12.1	4.10	9.8	22.3	3.38	4.06
	32	12.7	3.91	10.4	22.8	3.00	5.00
Pubmed	1	7.31	4.16	5.62	6.70	6.79	-8.70
	4	8.80	4.92	7.11	5.21	7.02	-10.5
	32	9.07	3.91	7.38	-10.5	19.8	-29.9
Athletes	1	8.73	2.81	6.56	12.6	2.88	-6.18
	4	12.1	1.89	9.88	18.2	3.09	0.70
	32	12.0	1.89	9.88	15.5	4.49	-2.62
Company	1	16.8	2.43	14.6	17.9	1.75	0.13
	4	19.4	1.83	17.2	22.3	1.71	5.54
	32	18.2	1.86	16.1	20.1	3.34	2.87
Government	1	10.3	2.43	8.07	15.4	3.06	-2.98
	4	11.7	2.75	9.45	22.0	1.89	5.06
	32	10.7	2.67	8.47	23.6	2.53	7.11
New sites	1	16.5	3.70	14.2	11.3	3.80	-8.70
	4	16.6	3.70	14.3	9.70	5.60	-10.7
	32	15.2	3.50	12.9	4.00	6.90	-17.7
Politicians	1	15.7	2.52	13.6	18.5	1.60	1.98
	4	17.2	2.19	15.1	24.6	1.60	9.44
	32	17.9	2.19	15.8	24.8	1.40	9.72
Public Figures	1	10.5	3.48	8.25	17.5	3.21	-0.70
	4	9.51	4.09	7.29	24.3	2.97	7.69
	32	10.2	5.08	7.93	26.9	3.41	10.9
TV shows	1	18.3	1.61	16.2	19.3	1.89	2.63
	4	21.3	1.40	19.1	25.2	1.55	9.88
	32	21.8	1.14	19.7	24.4	1.55	8.82

Table 6: Varying the power $p \in \{1, 4, 32\}$: Training gain and computation gain of selected datasets with respect to DEEPWALK baseline. For loss-guided with Lscore₁ and Lscore₁₀. We keep the update frequency fixed, $F = 10$.

higher loss. A lower p allows for broader selection of walks into a round and higher p focuses the training more on highest loss score walks. In particular, $p = 32$ means that we essentially select walks with highest-loss examples whereas $p = 1$ means that we select walks for training proportionately to the loss values of their example(s). The loss power p selection should be dependent on F , because lower F allows for a broader selection of walks as well. It also needs to be fitted to the t' we use. Table 6 reports training and computation gains when we vary the loss power $p \in \{1, 4, 32\}$. This for the loss scoring functions Lscore₁ and Lscore₁₀. We can see that higher p values, 4 or 32, perform better overall than $p = 1$ and that the improvement are fairly robust to the particular choice of p .

C RESULTS FOR NODE2VEC BASELINE

Plots for the quality in the course of training with the NODE2VEC baseline for representative datasets are provided in Figure 7. Recall that the respective training and computation gains of the loss guided methods were reported in Table 2 and Table 3.

Figure 7: NODE2VEC baseline: Quality in the course of training for the Athletes, Company, and Politicians networks from the Facebook collection, the Pubmed and Cora citation networks, and the PPI network.

D LOSS BEHAVIOR

We observed that loss-guided walk selection improves performance of the different downstream tasks. To obtain insights on the behavior of the loss-guided versus the baseline methods we consider properties of the per-example loss values. Figure 8 provides side-by-side plots of these properties and plots of quality in the course of training. For our purposes here, we treat the graph edges as an approximate set of strong positive examples. These examples tend to be weighted higher (have larger $\kappa_{i,j}$ values) in the distribution generated from random walks. We consider two qualities of the distribution of the loss values $L_+(i, j)$ on these edges:

- The ratio of the average loss on edges to the *background* loss. We measure the background loss by the average loss $L_+(i, j)$ over 10^3 random non-edge pairs (i, j) . The background loss serves as a normalizing quantity for the average edge loss. This is because we observed that the scale of loss values shifts significantly during training, and in particular, the average edge and background loss values both decrease by orders of magnitude. We expect the ratio to be lower when the training is more effective. A lower ratio means a larger separation between edge loss (which we want to be low) and background loss.
- The ratio of the 90% quantile of edge loss values $L_+(i, j)$ to the average edge loss. The ratio is a measure of spread and indicates how well the training method balances its progress across the positive examples. A ratio that is closer to 1 means a smaller spread and a better balance.

Results on representative datasets are reported in Figure 8. We can see that across datasets and in the training regime before performance peaks, the loss-guided methods have a lower spread than the baseline DEEPWALK method: The ratio of the 90% percentile to the average loss on edges is uniformly lower. Moreover, the loss-guided method with Lscore₁₀ has a lower spread than Lscore₁. Overall, this

Figure 8: Patterns of per-example loss values: The ratio of edges loss versus background loss (left), the 90% quantile versus average edge loss (middle) and quality in the course of training (right). For loss-guided methods with loss scoring $Lscore_1$ and $Lscore_{10}$ and the baseline method DEEPWALK.

is consistent with what we expect with loss-guided training, where more training is directed to higher loss examples.

Interestingly, the baseline method DEEPWALK has a lower ratio, which corresponds to stronger separation from the background loss. The lower ratio starts early on and surprisingly perhaps, on some of the datasets the larger separation than the loss-guided methods persists also in regimes where DEEPWALK is outperformed by the

loss-guided methods (see for example the PPI and citation networks datasets).

These patterns showcase the advantage of loss-guided selections that are more geared to minimize spread rather than average loss. The average loss seems to indeed be effectively minimized by baseline methods, but on its own provides a limited measure of overall quality.