

SketchNE: Embedding Billion-Scale Networks Accurately in One Hour

Yuyang Xie, Yuxiao Dong, Jiezhong Qiu, Wenjian Yu, Xu Feng, Jie Tang

Abstract—We study large-scale network embedding with the goal of generating high-quality embeddings for networks with more than 1 billion vertices and 100 billion edges. Recent attempts LightNE and NetSMF propose to sparsify and factorize the (dense) NetMF matrix for embedding large networks, where NetMF is a theoretically-grounded network embedding method. However, there is a trade-off between their embeddings’ quality and scalability due to their expensive memory requirements, making embeddings less effective under real-world memory constants. Therefore, we present the SketchNE model, a scalable, effective, and memory-efficient network embedding solution developed for a single machine with CPU only. The main idea of SketchNE is to avoid the explicit construction and factorization of the NetMF matrix either sparsely or densely when producing the embeddings through the proposed sparse-sign randomized single-pass SVD algorithm. We conduct extensive experiments on nine datasets of various sizes for vertex classification and link prediction, demonstrating the consistent outperformance of SketchNE over state-of-the-art baselines in terms of both effectiveness and efficiency. SketchNE costs only **1.0 hours** to embed the Hyperlink2012 network with **3.5 billion** vertices and **225 billion** edges on a CPU-only single machine with embedding superiority (e.g., a **282%** relative HITS@10 gain over LightNE).

Index Terms—network embedding, network representation learning, randomized matrix factorization, memory-efficient.



1 INTRODUCTION

REPRESENTATION learning on graphs has recently provided a new paradigm for modeling real-world networks [1]. Learning structural representations for networks, i.e., network embedding, aims to map network entities into a latent space. The learned entity embeddings have been used to power various billion-scale online services, such as DeepWalk [2] in Alibaba [3], LINE [4] in LinkedIn [5], metapath2vec [6] and NetSMF [7] in Microsoft Academic [8], PinSage in Pinterest [9].

Take Facebook for example, it leverages the word2vec [10] based graph embedding system [11] to learn structural embeddings for its 3 billion user base. These embeddings are then consumed in various downstream applications. To maintain the quality of these embeddings, it is required to periodically embed such networks as its underlying structure consistently evolves, ideally as frequently as possible, e.g., every few hours in Alibaba [3]. However, according to our estimates, state-of-the-art (SOTA) graph embedding systems, i.e., GraphVite [12]—a DeepWalk [2] based system—and PyTorch-BigGraph [11]—would cost days if not weeks by using powerful CPU and GPU clusters to embed a network of 3B users.

Though skip-gram based embedding models, e.g., DeepWalk [2], LINE [4], and metapath2vec [6], have been widely adopted in large-scale solutions. They are still limited to handle billion-scale networks at speed, as discussed above. Recently, a theoretical study demonstrates that these models can be transformed as implicit factorization of a closed-form matrix [13]. Based on this discovery, the NetMF model is

proposed to explicitly construct and factorize the matrix that is implicitly factorized by DeepWalk, i.e., the NetMF matrix¹ $f^\circ(M)$, in which M can be approximated by L and R —the two matrices formed by the eigen-decomposition over the graph Laplacian of a given network—and $f^\circ(\cdot)$ is an element-wise logarithm function. Additionally, addressing the matrix form $f^\circ(LR)$ can benefit various machine learning scenarios, such as the attention mechanism in Transformer [14]— $\text{softmax}(\cdot)$, the linear layer with ReLU [15] activation— $\text{ReLU}(\cdot)$, and the kernel method [16].

Despite its outperformance over skip-gram based methods, it is computationally prohibitive for NetMF to handle million-scale networks as it needs to construct and factorize $f^\circ(M)$, which is an $n \times n$ dense matrix with n being the number of vertices. To address this, one recent attempt NetSMF [7] proposes to construct a sparse version of $f^\circ(M)$ by a graph spectral based sampling technique and then leverage sparse matrix factorization to produce vertex embeddings. More recently, LightNE [17] advances NetSMF by further reducing its sampling cost, utilizing other system-wise optimizations, and borrowing the spectral propagation strategy from ProNE [18]. In doing so, LightNE outperforms SOTA systems, including NetSMF, ProNE, GraphVite, and PyTorch-BigGraph, in terms of both computational cost and embedding effectiveness.

However, the performance of LightNE and NetSMF heavily relies on the number of samplings that directly corresponds to the memory cost, that is, more samplings make the sparse matrix more close to $f^\circ(M)$ and thus yield better embeddings, while consuming more memory. For example, to generate competitive embeddings for the OAG data [17] of 67M vertices and 895M edges, LightNE requires 1493GB memory space to have a sufficient number

• The authors are all with the Department of Computer Science and Technology, Tsinghua University, Beijing, 100084, China. E-mail: xyj18@mails.tsinghua.edu.cn, yuxiaod@tsinghua.edu.cn, qiuwj16@mails.tsinghua.edu.cn, yu-wj@tsinghua.edu.cn, fx17@mails.tsinghua.edu.cn, jietang@tsinghua.edu.cn.

1. The detailed NetMF matrix $f^\circ(M)$ can be found in Table 1.

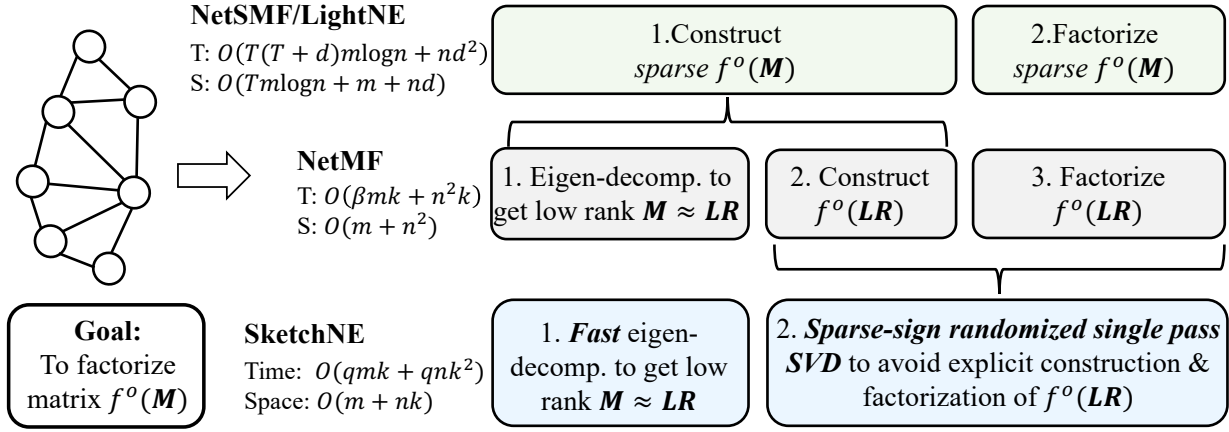


Fig. 1: The overview of SketchNE vs. NetMF and NetSMF/LightNE. The symbols used are listed in Table 1.

of samples. In order to embed larger networks, such as those of billions of vertices, LightNE has to sacrifice the quality of the embeddings under the real-world memory constraint.

Contributions. In light of the limitations of existing large-scale graph embedding solutions, the goal of this work is to learn effective embeddings for billion-scale networks efficiently under certain memory constraints, e.g., to embed networks with 3B vertices and 200B edges within 1 hour by using a single machine with 1500GB memory. To achieve this, we present the SketchNE² model, an effective, scalable, and memory-efficient method for billion-scale network embedding. Figure 1 illustrates the two technical components in SketchNE—a fast eigen-decomposition algorithm and a sparse-sign randomized single-pass SVD, each of which addresses the computational challenges in NetMF correspondingly.

First, we propose to factorize the target matrix $f^o(\mathbf{M})$ without explicitly constructing it, avoiding the direct or sparse construction and factorization in NetMF or NetSMF/LightNE. To achieve this, we present the sparse-sign randomized single-pass SVD algorithm by leveraging the concept of the randomized sketch matrix.

Second, the step above still requires \mathbf{L} and \mathbf{R} , though the explicit construction of $f^o(\mathbf{M})$ is not demanded anymore. Thus, we further introduce a fast randomized eigen-decomposition algorithm to approximate the computation of \mathbf{L} and \mathbf{R} and give an upper bound of the approximation error. Empirical tests show that we can achieve about 90× speedup over the original eigen-decomposition module in NetMF without performance loss on (small) networks that NetMF can actually handle.

Third, we combine the spectral propagation strategy which is proposed in [18] to further improve the quality of the initial embedding. We optimize our system for shared-memory architectures with Graph Based Benchmark Suite (GBBS) [19], which has already shown its superiority when handling real-world networks with hundreds of billions of edges on a single machine. Intel Math Kernel Library (MKL) is used in SketchNE for basic linear algebra operations.

2. The code is publicly available at <https://github.com/xyphant0m/SketchNE>

We conduct extensive experiments to examine the performance of SketchNE, including its effectiveness, efficiency, and memory cost. Specifically, we test SketchNE and other SOTA models/systems on five datasets for vertex classification and four datasets for link prediction. The results show that by using the least running time and memory among SOTA models/systems, SketchNE can consistently outperform nine large-scale baselines across five datasets for vertex classification and also offers significant improvements over LightNE on three billion-scale networks for link prediction. Notably, *SketchNE can embed the Hyperlink2012 network with 3.5 billion vertices and 225 billion edges in 1.0 hours by using 1,321GB memory on a single machine, and the learned embeddings offer a 282% relative HITS@10 improvement over LightNE on the link prediction task.*

2 NETMF AND ITS CHALLENGES

Given an undirected network $G = (V, E, \mathbf{A})$ with n vertices, m edges, adjacency matrix \mathbf{A} , degree matrix \mathbf{D} and volume $\text{vol}(G) = \sum_i \sum_j \mathbf{A}_{ij}$, the goal of network embedding is to learn an embedding matrix $\mathbf{E} \in \mathbb{R}^{n \times d}$ so that row i captures the structural property of vertex i [2], [4], [13], [20]. The embeddings can be then fed into downstream applications. The symbols are listed in Table 1.

Many network embedding algorithms are based on random walk and skip-gram techniques, such as DeepWalk [2], LINE [4], and node2vec [20]. Take DeepWalk for example, the vertex sequences traversed by random walkers are fed into the skip-gram model, which is usually parameterized by the context window size T and the number of negative samples b . Notably, these techniques are later shown to be

TABLE 1: Symbol used throughout this paper.

Symbol	Description	Symbol	Description
G	input network	b	#negative samples
V	vertex set, $ V = n$	T	context window size
E	edge set, $ E = m$	\mathbf{E}	embedding matrix
\mathbf{A}	adjacency matrix	d	embedding dimension
$\text{vol}(G)$	volume of G	\mathbf{U}_k	truncated eigenvectors
\mathbf{D}	degree matrix	Λ_k	truncated eigenvalues
$\mathbf{D}^{-\alpha} \mathbf{A} \mathbf{D}^{-\alpha}$	modified Laplacian	k	rank parameter
\mathcal{L}	normalized Laplacian	q	power parameter
\mathbf{M}	$\frac{\text{vol}(G)}{bT} \sum_{r=1}^T (\mathbf{D}^{-1} \mathbf{A})^r \mathbf{D}^{-1}$	\mathbf{Y}	sketch of $f^o(\mathbf{LR})$
s	oversampling parameter	z	column density

Algorithm 1: NetMF under the new formulation

Input: adjacency matrix \mathbf{A} , rank k , embedding dimension d
Output: An embedding matrix $\mathbf{E} \in \mathbb{R}^{n \times d}$

- 1 $[\mathbf{U}_k, \mathbf{\Lambda}_k] = \text{eigs}(\mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}, k)$
// Eigen-decomposition
- 2 $\mathbf{L} = \frac{\text{vol}(\mathbf{G})}{bT} \mathbf{D}^{-1/2} \mathbf{U}_k, \mathbf{R} = \left(\sum_{r=1}^T \mathbf{\Lambda}_k^r \right) \mathbf{U}_k^\top \mathbf{D}^{-1/2}$
- 3 $[\mathbf{U}_d, \mathbf{\Sigma}_d, \mathbf{V}_d] = \text{svds}(f^\circ(\mathbf{L}\mathbf{R}), d)$ // Rank- d truncated SVD
- 4 **return** $\mathbf{E} = \mathbf{U}_d \mathbf{\Sigma}_d^{1/2}$ as network embedding

theoretically equivalent to matrix factorization [13]. Based on this result, the NetMF algorithm is proposed to explicitly construct and factorize the matrix that is implicitly factorized by DeepWalk, namely the NetMF matrix:

$$\text{trunc_log}^\circ \left(\frac{\text{vol}(\mathbf{G})}{bT} \sum_{r=1}^T (\mathbf{D}^{-1} \mathbf{A})^r \mathbf{D}^{-1} \right), \quad (1)$$

where trunc_log° denotes the element-wise truncated logarithm, i.e., applying $\text{trunc_log}(x) = \max(0, \log(x))$ to each entry of a matrix. However, the explicit construction and factorization of this matrix usually consumes $O(n^3)$ time as it tends to be a dense matrix even with a small T . To reduce time complexity, NetMF conducts truncated eigen-decomposition such that $\mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2} \approx \mathbf{U}_k \mathbf{\Lambda}_k \mathbf{U}_k^\top$, and factorizes the following approximate matrix of (1):

$$\text{trunc_log}^\circ \left(\frac{\text{vol}(\mathbf{G})}{bT} \mathbf{D}^{-1/2} \mathbf{U}_k \left(\sum_{r=1}^T \mathbf{\Lambda}_k^r \right) \mathbf{U}_k^\top \mathbf{D}^{-1/2} \right). \quad (2)$$

Reformulate the Goal of NetMF. With the above description, we can reformulate and generalize the goal of NetMF as follows:

Problem 1. *Truncated SVD for element-wise function of a low-rank matrix.*

Given: Two low-rank matrices $\mathbf{L} \in \mathbb{R}^{n \times k}$ and $\mathbf{R} \in \mathbb{R}^{k \times n}$, a function $f : \mathbb{R} \rightarrow \mathbb{R}$ applied to each entry of $\mathbf{L}\mathbf{R}$, and desired dimensionality d .

Goal: Compute the rank- d truncated SVD for $f^\circ(\mathbf{L}\mathbf{R})$ such that:

$$[\mathbf{U}_d, \mathbf{\Sigma}_d, \mathbf{V}_d] = \text{svds}(f^\circ(\mathbf{L}\mathbf{R}), d). \quad (3)$$

In NetMF, $\mathbf{L} = \frac{\text{vol}(\mathbf{G})}{bT} \mathbf{D}^{-1/2} \mathbf{U}_k, \mathbf{R} = \left(\sum_{r=1}^T \mathbf{\Lambda}_k^r \right) \mathbf{U}_k^\top \mathbf{D}^{-1/2}$ and $f(\cdot) = \text{trunc_log}(\cdot)$.

Alg. 1 describes NetMF under the above new formulation. Unfortunately, it is still not capable of handling large networks, even the YouTube dataset with 1.1 million vertices used in DeepWalk [2], [21], mainly due to the following two challenges presented in Problem 1.

Challenge One: How to Solve $\text{svds}(f^\circ(\mathbf{L}\mathbf{R}), d)$ Efficiently?

The major challenge lies in the requirement to explicitly construct and factorize $f^\circ(\mathbf{L}\mathbf{R})$, even after NetMF’s attempt to perform the truncated eigen-decomposition. In fact, its construction and factorization in Alg. 1 Line 3 demand $O(n^2)$ memory cost and $O(n^2k)$ time cost, making it computationally infeasible for large networks. It is worth noting that the element-wise truncated logarithm is very important to embedding quality and cannot be omitted [7], [13], [17], [22]. Otherwise, the embeddings can be realized without

Algorithm 2: The basic randomized SVD

- 1 **Procedure** basic_randomized_SVD(\mathbf{X}, k, q)
- 2 $\mathbf{\Omega} = \text{randn}(n, k+s)$ // oversampling
parameter $s: 10$ or 20
- 3 $\mathbf{Y} = \mathbf{X}\mathbf{\Omega}$ // sketch matrix $\mathbb{R}^{n \times (k+s)}$
- 4 $\mathbf{Q} = \text{orth}(\mathbf{Y})$ // QR factorization
- 5 **for** $i = 1, \dots, q$ **do** // power iteration (optional)
- 6 | $\mathbf{T} = \text{orth}(\mathbf{X}^\top \mathbf{Q})$ // QR factorization
- 7 | $\mathbf{Q} = \text{orth}(\mathbf{X}\mathbf{T})$ // QR factorization
- 8 $\mathbf{B} = \mathbf{Q}^\top \mathbf{X}$ // reduced matrix $\mathbf{B} \in \mathbb{R}^{(k+s) \times n}$
- 9 $[\hat{\mathbf{U}}, \hat{\mathbf{\Sigma}}, \hat{\mathbf{V}}] = \text{svd}(\mathbf{B})$ // Full SVD
- 10 $\mathbf{U} = \mathbf{Q}\hat{\mathbf{U}}(:, 1:k), \mathbf{\Sigma} = \hat{\mathbf{\Sigma}}(1:k, 1:k), \mathbf{V} = \hat{\mathbf{V}}(:, 1:k)$
- 11 **return** $\mathbf{U}, \mathbf{\Sigma}, \mathbf{V}$

constructing the dense form, as in NRP [23], RandNE [24], and FastRP [25].

Challenge Two: How to Factorize $\mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}$ Efficiently? Although one may think the truncated eigen-decomposition of $\mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}$ (Line 1 of Alg. 1) is a simple and efficient step, previous work [26], [27] and our analysis show that it is in fact computationally very expensive. In particular, Cohen-Steiner et al. [26] shows it is very difficult to obtain the spectrum of a large graph, and the eigen-decomposition of a large graph Laplacian [27] is practically very slow. The cost of the truncated eigen-decomposition [28] is $O(\beta mk)$, where $\beta \geq 1$ and its value depends on the convergence speed. The convergence, in turn, depends on the relative eigenvalue gap, making this constant term very big and this operation practically very expensive. This problem seems not prominent in NetMF because the datasets in its paper [13] are relatively small (the largest one is Flickr with 80K vertices). However, for a slightly larger dataset YouTube with 1.1M vertices, we observe that the SciPy implementation `eighsh` cannot finish the computation in three days, not to mention for billion-scale networks.

3 THE SKETCHNE MODEL

In this section, we present SketchNE for embedding billion-scale networks at speed. In Sec. 3.1, we propose a sparse-sign single-pass SVD algorithm to resolve challenge one—factorizing $f^\circ(\mathbf{L}\mathbf{R})$ without constructing its dense form. In Sec. 3.2, we introduce a fast randomized eigen-decomposition algorithm to resolve challenge two—accelerating the computation of \mathbf{L} and \mathbf{R} . Sec. 3.3 combines the two sketching techniques and presents the overall algorithm.

3.1 Sketching-Based $\text{svds}(f^\circ(\mathbf{L}\mathbf{R}), d)$ without Explicitly Computing & Factorizing $f^\circ(\mathbf{L}\mathbf{R})$

In this part, we formally introduce a sketch-based solution to $\text{svds}(f^\circ(\mathbf{L}\mathbf{R}), d)$ without explicitly computing and factorizing it.

Basic Randomized SVD. To solve the computational challenge of $\text{svds}(f^\circ(\mathbf{L}\mathbf{R}), d)$, we first revisit the randomized SVD in Alg. 2, where $\mathbf{\Omega}$ (Line 2) is a Gaussian random

matrix, s (Line 2) is the oversampling parameter for better accuracy, q (Line 5) is the power iteration index, and $\text{orth}(\cdot)$ (Lines 4, 6 and 7) is the orthonormalization operation which is usually implemented with QR factorization. The random projection in Line 3 generates a sketch matrix Y , which identifies a subspace that captures dominant information of the input matrix $X = f^\circ(LR)$. Then, with the subspace's orthonormal basis matrix Q computed in Line 4, one obtains the approximation $X \approx QQ^\top X$ [29]. Lines 5-8 are the power iteration scheme which is optional to improve the approximation accuracy. Next, Line 9 constructs a reduced matrix B by projecting the input matrix X to the subspace with orthonormal basis Q . Finally, the approximate truncated SVD of X is obtained through performing SVD on matrix B in Lines 10-11.

We can see that there are four places in Alg. 2 (Lines 3, 6, 7 and 9) requiring the explicit construction of $X = f^\circ(LR)$. For Lines 6-7, we can avoid them by skipping the optional power iteration. We discuss how to deal with issues raised by Lines 3 and 9, respectively.

Issue One (Alg. 2 Line 3): Sketching $f^\circ(LR)$. Computing sketch matrix Y requires matrix multiplication between $X = f^\circ(LR)$ and a random matrix Ω . However, $f^\circ(LR)$ cannot be explicitly computed due to its $O(n^2k)$ time complexity and $O(n^2)$ memory cost.

Solution: The Sparse-Sign Matrix for Sketching. We introduce the concept of the sparse-sign matrix to help solve the multiplication $X\Omega = f^\circ(LR)\Omega$. The sparse-sign matrix, with similar performance to the Gaussian matrix, is another type of randomized dimension reduction map [30]. Alg. 3 describes how to generate a sparse-sign matrix. To have a sparse-sign matrix $S \in \mathbb{R}^{n \times h}$ with sketch size $h = d + s$, where d is the embedding dimension, and s is the oversampling parameter, we fix a column density parameter z in the range $2 \leq z \leq n$. We independently generate the columns of the matrix at random. For each column, we draw z i.i.d random signs and place them in z uniformly random coordinates as shown in Alg. 3 Lines 5-6. Line 8 is to get the unique row coordinates. According to [30], $z = \min(n, 8)$ will usually be a good choice.

After generating the sparse-sign matrix S , we can use it as the random matrix, which multiplies $f^\circ(LR)$ on its right side to obtain the sketch matrix Y . Considering that a column of the sparse-sign matrix S only has z nonzeros, it will generate at most $z \times h$ coordinates in the range of $[1, n]$ and $z \times h \ll n$. Therefore, we can perform column sampling according to these unique coordinates p . Assuming that $v = \text{size}(p, 1)$, then v satisfy $v \leq z \times h \ll n$. Based on the fact that S has only v non-zero rows, we can immediately observe that we can have a sampling matrix $R(:, p)$. Therefore, computing $Y = f^\circ(LR)S$ is exactly equivalent to

$$Y = f^\circ(LR(:, p))S. \quad (4)$$

The time cost of (4) is $O(nvk + nvh)$ and the memory cost is $O(nk + nv)$. However, when calculating $LR(:, p)$ for a network with billions of vertices, it will introduce $O(nv)$ memory cost, which is still infeasible. To solve this, we adopt the batch matrix multiplication by selecting the fixed-size rows of L in turn to complete the multiplication, which further reduces the memory cost to $O(nk)$.

Algorithm 3: Generate a sparse-sign matrix

```

1 Procedure gen_sparse_sign_matrix( $n, h, z$ )
2    $S = \text{zeros}(n, h)$  // sketch size  $h$ 
3    $p = \text{zeros}(zh, 1)$  // column density
   parameter  $z$ 
4   for  $j = 1, 2, \dots, h$  do
5      $p((z(j-1)+1) : zj) = \text{randperm}(n, z)$ 
6      $S(p((z(j-1)+1) : zj), j) = \text{sign}(\text{randn}(z, 1))$ 
7    $p = \text{unique}(p)$  // make coordinates unique
8   return  $S, p$ 

```

Issue Two (Alg. 2 Line 9): Form the Reduced Matrix B . According to Alg. 2 Line 9, the reduced matrix B is constructed by $B = Q^\top f^\circ(LR)$, where Q is a dense matrix and $f^\circ(LR)$ is implicitly stored, making it too expensive to obtain B .

Solution: The Randomized Single-Pass SVD. We leverage the idea of randomized single-pass SVD [30] to solve this issue. The basic idea is to obtain the approximate SVD results by visiting the target matrix $f^\circ(LR)$ only once. The process of single-pass SVD is as follows: First, we draw multiple sketch matrices that capture the row and column dominant information of matrix $f^\circ(LR)$ and compute SVD based on these sketch matrices. In [30], four sparse-sign random matrices $C \in \mathbb{R}^{n \times h}$, $S \in \mathbb{R}^{n \times h}$, $H \in \mathbb{R}^{n \times l}$, $O \in \mathbb{R}^{n \times l}$ are drawn for target matrix $f^\circ(LR)$. Then three sketch matrices

$$K = f^\circ(LR)^\top C, Y = f^\circ(LR)S, Z = H^\top f^\circ(LR)O \quad (5)$$

are generated respectively. Second, we obtain the orthonormal matrices

$$P = \text{orth}(K), Q = \text{orth}(Y), \quad (6)$$

which capture the row and column dominant information of $f^\circ(LR)$, respectively. Then we get a great approximation as

$$f^\circ(LR) \approx QQ^\top f^\circ(LR)PP^\top. \quad (7)$$

By updating Z with Eq. 7, we can have

$$Z = H^\top f^\circ(LR)O \approx (H^\top Q)(Q^\top f^\circ(LR)P)(P^\top O). \quad (8)$$

Third, we get the reduced matrix

$$W = (H^\top Q)^\dagger Z (P^\top O)^\dagger \approx Q^\top f^\circ(LR)P \quad (9)$$

by solving the least-squares problem. Finally, it will form a low rank approximation of the target matrix $f^\circ(LR)$ via

$$f^\circ(LR) \approx QWP^\top, \quad (10)$$

and the approximate truncated SVD of $f^\circ(LR)$ can be derived from performing SVD on W . Therefore, $f^\circ(LR)$ is required only in the sketching process, and the reduced matrix W is constructed only by sketch matrices and random matrices.

In addition, we note that $f^\circ(LR)$ is symmetric in NetMF, thus the row dominant information is equal to column dominant information which means we can omit K and H . This enables us to replace P with Q in Eq. 7-10 and replace H with O in Eqs. 5, 8, and 9. In other words, we can further simplify and improve the above randomized single-pass SVD process: When the multiplication $Y = f^\circ(LR)S$ is performed, we can simultaneously draw another sketch matrix $Z = O^\top f^\circ(LR)O$ with a sparse-sign random ma-

Algorithm 4: Sparse-sign randomized single-pass SVD

```

1 Procedure sparse_sign_rand_single_pass_SVD ( $\mathbf{L}, \mathbf{R}, d$ )
2    $h = d + s_1, l = d + s_2$  // sketch size  $h, l$ 
3    $[\mathbf{S}, \mathbf{p}] = \text{gen\_sparse\_sign\_matrix}(n, h, z)$  // Alg. 3
4    $\mathbf{Y} = f^\circ(\mathbf{L}\mathbf{R}(:, \mathbf{p}))\mathbf{S}$  // sketch matrix  $\mathbf{Y} \in \mathbb{R}^{n \times h}$ 
5    $\mathbf{Q} = \text{orth}(\mathbf{Y})$ 
6    $[\mathbf{O}, \mathbf{p}] = \text{gen\_sparse\_sign\_matrix}(n, l, z)$  // Alg. 3
7    $\mathbf{Z} = \mathbf{O}^\top f^\circ(\mathbf{L}(\mathbf{p}, :)\mathbf{R}(:, \mathbf{p}))\mathbf{O}$  // sketch matrix
    $\mathbf{Z} \in \mathbb{R}^{l \times l}$ 
8    $\mathbf{W} = (\mathbf{O}^\top \mathbf{Q})^\dagger \mathbf{Z} (\mathbf{Q}^\top \mathbf{O})^\dagger$  // reduced matrix
    $\mathbf{W} \in \mathbb{R}^{h \times h}$ 
9    $[\hat{\mathbf{U}}, \hat{\Sigma}, \hat{\mathbf{V}}] = \text{svd}(\mathbf{W})$  // Full SVD
10   $\mathbf{U} = \mathbf{Q}\hat{\mathbf{U}}(:, 1:d), \Sigma = \hat{\Sigma}(1:d, 1:d), \mathbf{V} = \mathbf{Q}\hat{\mathbf{V}}(:, 1:d)$ 
11  return  $\mathbf{U}, \Sigma, \mathbf{V}$ 
    
```

trix $\mathbf{O} \in \mathbb{R}^{n \times l}$.

Overall: Sparse-Sign Randomized Single-Pass SVD. By combining the sparse-sign random matrix with single-pass SVD, we propose a sparse-sign randomized single-pass SVD algorithm to avoid the explicit construction and factorization of $f^\circ(\mathbf{L}\mathbf{R})$ as Alg. 4.

In Alg. 4, Line 8 generates the reduced matrix \mathbf{W} , which involves solving the least-squares problem twice. The first is to solve $(\mathbf{O}^\top \mathbf{Q})\mathbf{T} = \mathbf{Z}$ for the temporary matrix \mathbf{T} and the second is to solve $(\mathbf{O}^\top \mathbf{Q})\mathbf{W}^\top = \mathbf{T}^\top$ for the reduced matrix \mathbf{W} . The matrix $\mathbf{O}^\top \mathbf{Q} \in \mathbb{R}^{l \times h}$ is well-conditioned when $l \gg h$ which suggests choosing $s_2 \gg s_1$. Lines 3 and 6 generate two sparse-sign random matrices. Considering that s_1 and s_2 are small numbers, the time cost of Line 4 is $O(nv(k+d))$ and Line 7 costs $O(v^2(k+d) + vd^2)$. Line 5 costs $O(nd^2)$, Lines 8-9 cost $O(d^3)$, and Line 10 costs $O(nd^2)$. Therefore, the time complexity of Alg. 4 is $O(n(vk + vd + d^2))$, linear to the number of vertices n , which is much more efficient than the explicit construction and factorization in NetMF ($O(n^2k)$).

Overall, we present the sparse-sign randomized single-pass SVD to address the computational challenges in NetMF, which is the first attempt to introduce single-pass low-rank matrix factorization into network embedding. It not only solves the challenges of NetMF, but also gives a solution to the general problem of factorizing $f^\circ(\mathbf{L}\mathbf{R})$. Recently, Han et al. [31] proposes polynomial tensor sketch for this problem, which combines a polynomial approximation of f (e.g., Taylor and Chebyshev expansion) with tensor sketch for approximating monomials of entries of $\mathbf{L}\mathbf{R}$. We will see an ablation study which applies polynomial tensor sketch to NetMF in Sec. 4.4.

3.2 Fast Eigen-Decomposition via Sketching

By now, we bypass the major bottleneck of Alg. 1 (Lines 3) without explicitly computing $f^\circ(\mathbf{L}\mathbf{R})$. However, the solution in Alg. 4 still requires the separate \mathbf{L} and \mathbf{R} as input, which are computed by the eigen-decomposition on $\mathbf{D}^{-1/2}\mathbf{A}\mathbf{D}^{-1/2}$ in Alg. 1 Line 1. Though the truncated eigen-decomposition costs only $O(\beta mk)$ FLOPs in theory ($\beta \geq 1$) [28], it is in practice almost infeasible to handle large networks due to the big constants in its complexity. In fact, the computation of this step for the YouTube dataset with 1.1 million vertices cannot complete within three days by using the commonly-used eigsh implementation, while the goal

Algorithm 5: Fast randomized eigen-decomposition

```

1 Procedure freigs ( $\mathbf{X}, k, q$ )
2    $[\sim, n] = \text{size}(\mathbf{X}, 2)$ 
3    $\Omega = \text{randn}(n, k + s)$ 
4    $\mathbf{Y} = \mathbf{X}\Omega$  // sketch matrix  $\mathbf{Y} \in \mathbb{R}^{n \times (k+s)}$ 
5    $[\mathbf{Q}, \sim, \sim] = \text{eigSVD}(\mathbf{Y})$  // fast
   orthonormalization
6   for  $i = 1, 2, \dots, q$  do // power iteration
7      $[\mathbf{Q}, \sim, \sim] = \text{eigSVD}(\mathbf{X}\mathbf{X}\mathbf{Q})$ 
8      $\mathbf{S} = \mathbf{Q}^\top \mathbf{X}\mathbf{Q}$  // reduced matrix  $\mathbf{S} \in \mathbb{R}^{(k+s) \times (k+s)}$ 
9      $[\hat{\mathbf{U}}, \Lambda] = \text{eig}(\mathbf{S})$  // Full eigen-decomposition
10     $\mathbf{U} = \mathbf{Q}\hat{\mathbf{U}}$ 
11    return  $\mathbf{U}(:, 1:k), \Lambda(1:k, 1:k)$ 
    
```

of this work is to embed billion-scale networks efficiently, e.g., in one hour.

To address this practical challenge, we introduce a fast randomized eigen-decomposition method to approximate $\mathbf{D}^{-1/2}\mathbf{A}\mathbf{D}^{-1/2}$. According to [29], the symmetric approximation formula should be $\mathbf{X} \approx \mathbf{Q}\mathbf{Q}^\top \mathbf{X}\mathbf{Q}\mathbf{Q}^\top$ and the truncated eigen-decomposition result of \mathbf{X} can be derived by performing eigen-decomposition on the small matrix $\mathbf{Q}^\top \mathbf{X}\mathbf{Q}$. By combining the techniques of the power iteration scheme and acceleration strategy [32], the fast randomized eigen-decomposition can be described as Alg. 5.

Practically, a good decomposition of $\mathbf{D}^{-1/2}\mathbf{A}\mathbf{D}^{-1/2}$ by freigs requires a large q , increasing the time cost (see the experiment in Sec. 4.4). To balance the trade-off between effectiveness and efficiency, we propose to perform Alg. 5 on a modified Laplacian matrix $\mathbf{D}^{-\alpha}\mathbf{A}\mathbf{D}^{-\alpha}$, where $\alpha \in (0, 0.5]$. Therefore, we have $\mathbf{U}_k \Lambda_k \mathbf{U}_k^\top \approx \mathbf{D}^{-\alpha}\mathbf{A}\mathbf{D}^{-\alpha}$. It means $\mathbf{D}^{-1/2}\mathbf{A}\mathbf{D}^{-1/2}$ is computed as $\mathbf{D}^{-1/2+\alpha}\mathbf{U}_k \Lambda_k \mathbf{U}_k^\top \mathbf{D}^{-1/2+\alpha}$ approximately. We give an upper bound of the approximation error by Lemma 4 and its proof both in Appendix A. In doing so, Eq. (1) can be approximated by

$$f^\circ(\mathbf{L}'\mathbf{R}') = f^\circ\left(\frac{\text{vol}(G)}{bT}\mathbf{D}^{-1+\alpha}\mathbf{U}_k \Lambda_k \left(\sum_{r=1}^T \mathbf{K}^{r-1}\right) \mathbf{U}_k^\top \mathbf{D}^{-1+\alpha}\right), \quad (11)$$

where $\mathbf{L}' = \frac{\text{vol}(G)}{bT}\mathbf{D}^{-1+\alpha}\mathbf{U}_k$, $\mathbf{K} = \mathbf{U}_k^\top \mathbf{D}^{-1+2\alpha}\mathbf{U}_k \Lambda_k$, $\mathbf{R}' = \Lambda_k (\sum_{r=1}^T \mathbf{K}^{r-1}) \mathbf{U}_k^\top \mathbf{D}^{-1+\alpha}$ and $f(\cdot) = \text{trunc_log}(\cdot)$. \mathbf{K} is a $k \times k$ matrix, making the computation of $\sum_{r=1}^T \mathbf{K}^{r-1}$ cheap. We further give an upper bound of the approximation error between the NetMF matrix and $f^\circ(\mathbf{L}'\mathbf{R}')$ by the following theorem. We can see the approximation is better with a larger k .

Theorem 1. Suppose f° denotes trunc_log° , i.e. the element-wise truncated logarithm, $f^\circ(\mathbf{M})$ is the matrix in (1), and $f^\circ(\mathbf{L}'\mathbf{R}')$ is defined by (11) which includes the quantities obtained with Alg. 5. Then,

$$\|f^\circ(\mathbf{M}) - f^\circ(\mathbf{L}'\mathbf{R}')\|_F \leq \frac{(1 + \varepsilon)d_{\min}^{-1+2\alpha} B}{(c - 1)bT}$$

with high probability. Here $|\lambda_j|$ is the j -th largest absolute value of eigenvalue of $\mathbf{D}^{-\alpha}\mathbf{A}\mathbf{D}^{-\alpha}$, $B = \sqrt{\sum_{j=k+1}^n |\lambda_j|^2 ((c^T - 1)(1 + \frac{n}{c-1}) - nT)\text{vol}(G)}$, $c = n(\frac{d_{\max}}{d_{\min}})^{1-2\alpha}$. d_{\min} and d_{\max} are the minimum and maximum vertex degrees, respectively.

Proof. See Appendix A. □

Algorithm 6: SketchNE

Input: A network $G = (V, E, \mathbf{A})$; Normalized parameter α ;
rank parameter k ; power parameter q ; Embedding
dimension d ;
Output: An embedding matrix $\mathbf{E} \in \mathbb{R}^{n \times d}$

- 1 $[\mathbf{U}_k, \mathbf{\Lambda}_k] = \text{freigs}(\mathbf{D}^{-\alpha} \mathbf{A} \mathbf{D}^{-\alpha}, k, q)$ // Alg. 5
- 2 $\mathbf{K} = \mathbf{U}_k^\top \mathbf{D}^{-1+2\alpha} \mathbf{U}_k \mathbf{\Lambda}_k$
- 3 $\mathbf{L}' = \frac{\text{vol}(G)}{bT} \mathbf{D}^{-1+\alpha} \mathbf{U}_k, \mathbf{R}' = \mathbf{\Lambda}_k \left(\sum_{r=1}^T \mathbf{K}^{r-1} \right) \mathbf{U}_k^\top \mathbf{D}^{-1+\alpha}$
- 4 $[\mathbf{U}_d, \mathbf{\Sigma}_d, \sim] = \text{sparse_sign_rand_single_pass_SVD}(\mathbf{L}', \mathbf{R}', d)$
// Alg. 4
- 5 $\mathbf{E} = \mathbf{U}_d \mathbf{\Sigma}_d^{1/2}$ // inital embedding
- 6 $\mathbf{E} = \sum_{r=0}^p c_r \mathcal{L}^r \mathbf{E}$ // spectral propagation
- 7 **return** \mathbf{E} as network embedding

In Alg. 5, the “eigSVD” is used as the orthonormalization operation. Compared with the QR factorization, eigSVD is much faster especially when $n \gg (k + s)$ [32]. Since the oversampling parameter s is smaller than k , Lines 4-5 cost $O(mk + nk^2)$. According to [32], Lines 6-8 cost $O(q(mk + nk^2))$, Line 9 costs $O(mk + nk^2)$, Line 10 costs $O(k^3)$, and Line 11 costs $O(nk^2)$. Overall, the time complexity of Alg. 5 is $O(q(mk + nk^2))$. However, the actual FLOPs of Alg. 5 is far fewer than that of eigsh. In practice, our empirical tests suggest that by setting $\alpha=0.4$, Alg. 5 with a small q shows on average -90X speedup to eigsh on the small datasets that can be handled by eigsh without noticeable impacts on the eigenvalues computed and by extension on the embeddings learned (See Fig. 3 (a)).

3.3 The Overall Algorithm

In Sec. 3.1, we develop a sparse-sign randomized single-pass SVD algorithm (Alg. 4) to solve $\text{svds}(f^\circ(\mathbf{LR}), d)$ without the explicit computation and factorization of the full matrix $f^\circ(\mathbf{LR})$. In Sec. 3.2, we propose a fast randomized eigen-decomposition method to get \mathbf{L} and \mathbf{R} for large networks. Empowered by these two techniques, we address the two computational challenges faced by NetMF, respectively. To this end, we present the SketchNE algorithm to learn embeddings for billion-scale networks in Alg. 6.

In Alg. 6, Line 1 computes the fast randomized eigen-decomposition for the modified Laplacian matrix $\mathbf{D}^{-\alpha} \mathbf{A} \mathbf{D}^{-\alpha}$. Lines 2-3 form the approximations of matrices \mathbf{L} and \mathbf{R} . Line 4 is to compute the SVD of $f^\circ(\mathbf{L}'\mathbf{R}')$ by \mathbf{L}' and \mathbf{R}' through the sparse-sign randomized single-pass SVD. Then we form the inital embedding in Line 5. Line 6 conducts spectral propagation, which is a commonly-used and computationally-cheap enhancement technique [17], [18] that further improves embedding quality. c_r in Line 6 is the coefficients of Chebyshev polynomials, p is spectral propagation steps (the default setting is 10) and $\mathcal{L} = \mathbf{I} - \mathbf{D}^{-1} \mathbf{A}$ is normalized graph Laplacian matrix (\mathbf{I} is the identity matrix).

Complexity of SketchNE. For Line 1, the input matrix $\mathbf{D}^{-\alpha} \mathbf{A} \mathbf{D}^{-\alpha}$ is still an $n \times n$ sparse matrix and has $2m$ nonzeros. According to Section 3.2, it requires $O(q(mk + nk^2))$ time and $O(m + nk)$ space. As for Lines 2-3, $O(nk^2)$ time and $O(nk)$ space are required. The time cost of Line 4 is $O(n(d^2 + vk + vd))$ and its space cost is $O(n(k + d))$. Lines 5-6 demand $O(pmd + nd)$ time and $O(m + nd)$ space. In total, the SketchNE has the time complexity of

$O(q(mk + nk^2) + nd^2 + nvk + nvd + pmd)$ and the space complexity of $O(m + nk)$. Therefore, there is a trade-off between efficiency and effectiveness on the choice of q . In practice, we can easily find q that offers clear superiority on both efficacy and efficiency, including both memory cost and computing time, over existing large-scale network embedding techniques. Consider that q, k, d, v, p are all very small compared to m and n , the overall time complexity of SketchNE is linear to the number of edges m and the number of vertices n .

3.4 Implementation Details

Memory Reduction. Considering the memory cost of SketchNE is $O(m + nk)$, while NetSMF/LightNE ties performance to memory cost. Therefore, we consider to further optimize the memory cost of SketchNE by the Graph Based Benchmark Suite (GBBS) [19], which is an extension of the Ligra [33] interface. We optimize the memory cost of SketchNE with the Graph Based Benchmark Suite (GBBS) [19], which is an extension of the Ligra [33] interface. The GBBS is easy to use and has already shown its practicality for many large scale fundamental graph problems. LightNE [17] introduces GBBS to network embedding problems and shows its superiority to real-world networks with hundreds of billions of edges. The main benefit of GBBS to SketchNE is the data compression. A sparse adjacency matrix is usually stored in the compressed sparse row (CSR) format, which is also regarded as an excellent compressed graph representation [34]. However, the CSR format still incurs a huge memory overhead for the networks with hundreds of billions of edges. For example, storing a network with 1 billion vertices and 100 billion edges costs 1121 GB memory. Therefore, we need to compress it further and reduce memory cost. The GBBS can be regarded as a compressed CSR format for the graph from Ligra+ [35], which supports fast parallel graph encoding and decoding.

Parallelization. Two major computational steps of the SketchNE are sparse matrix-matrix multiplication (SPMM) and matrix-matrix product (GEMM), which are well supported by the Intel MKL library. SPMM operation is well supported by MKL’s Sparse BLAS Routine. However, MKL’s Sparse BLAS Routine requires the sparse matrix in CSR format as the input, which contradicts the original intention of using GBBS. Fortunately, GBBS supports traversing all neighbors of a vertex u for the compressed CSR format, and we can propose an SPMM operation with the help of GBBS. In order to use GBBS to save memory cost, we propose a parallel GBBS-based SPMM operation to replace the SPMM operation in MKL’s sparse BLAS routine. The parallel GBBS-based SPMM is implemented as follows. Firstly, we traverse n vertices parallelly. Then, we traverse neighbor vertex v of vertex u to compute the quantity $\mathbf{D}(u, u)^{-\alpha} \mathbf{D}(v, v)^{-\alpha}$ corresponding to the sparse matrix. Finally, with the support of `cblas_saxpy` in MKL, we multiply the v -th row of the row-major matrix with the quantity and add the result to the u -th row of the target matrix. The SPMM operation based on GBBS is slightly slower than MKL’s SPMM operation, but ensuring memory-efficient. The “eigSVD”, “eig” and other operations in SketchNE are well supported by Intel

MKL routines. Line 4 and Line 7 of the Alg. 6 involve matrix column sampling and batch GEMM operation. They are easily parallelized with the “parallel for” derivative in OpenMP [36].

In conclusion, SketchNE is implemented in C++. We use GBBS to reduce memory usage and implement a GBBS-based SPM operation. For better efficiency, we use the Intel MKL library for basic linear algebra operations and the OpenMP programming.

4 EXPERIMENTS

In this section, we evaluate SketchNE on multi-label vertex classification and link prediction tasks, following exactly the same experimental settings as existing studies [2], [4], [7], [11], [12], [13], [17], [20]. We introduce datasets in Section 4.1 our experimental settings and results in Section 4.2 and Section 4.3, respectively. The ablation and case studies is in Section 4.4.

4.1 Datasets

We employ five datasets for the multi-label vertex classification task. BlogCatalog and YouTube are small graphs with less than 10 million edges, while the others are large graphs with more than 10 million but less than 10 billion edges. For the link prediction task, We have four datasets in which vertex labels are not available. Livejournal is the large graph, while the others are very large graphs with more than 10B edges. These datasets are of different scales and but have been widely used in network embedding literature [7], [17]. The statistics of datasets are listed in Table 2.

BlogCatalog [37] is a network of relationships of online users. The labels represent the interests of the users.

YouTube [21] is a video-sharing website, which allows user to upload, view, rate and share videos. The vertex labels represent the user’s taste in the video.

Friendster-small [38] is a sub-graph induced by all the labeled vertices in Friendster. The vertex labels in this network are the same as those in Friendster.

Friendster [38] is a large social network in an online gaming site. For some of the vertices, they have labels representing the groups the user joined.

OAG [39] is a publicly available academic graph opened by Microsoft Academic [39] and AMiner.org [40]. The vertex labels represent the study fields of each author.

Livejournal [41] is an online blogging site, where users can follow others to form a large social network.

ClueWeb [42] was created to support research on information retrieval and related human language technologies. The links between webs form the very large graphs.

Hyperlink2014 [43] was extracted from the Common Crawl Corpus released in April 2014, covering 1.7 billion web pages and 124 billion hyperlinks between these pages.

Hyperlink2012 [43] was extracted from the 2012 Common Crawl Corpus covering 3.5 billion web pages and 225 billion hyperlinks between these pages.

4.2 Experimental Settings

Baselines and Hyper-parameters Setting. We compare SketchNE with nine SOTA network embedding methods, including PyTorch-BigGraph (PBG) [11], GraphVite [12], NetMF [13], NetSMF [7], ProNE [18], LightNE [17], RandNE [24], FastRP [25] and NRP [23]. We also compare SketchNE with two GNN methods, including DGI [44] and GraphSAGE [45]. For all the baselines originally run on CPU and SketchNE, we test them with all the datasets with 88 threads on a server with two Intel[®] Xeon[®] E5-2699 v4 CPUs (88 virtual cores in total) and 1.5 TB memory. For GraphVite, we present the results obtained from the original paper (if existed), which uses a 4×P100 GPU server, and otherwise run it on a 4×V100 GPU server to get the results. For GNN methods, we run it on a server with 8× GeForce RTX 2080 Ti GPUs to get the results. All the baselines are evaluated with the hyperparameters set default in the corresponding paper’s GitHub Repository or tuned for the best performance. Their settings are as follows.

NetMF [13]. We download the authors’ official source codes³, and run experiments with default setting: $T = 10$, $k = 256$.

RandNE [24]. We download the authors’ official source codes⁴, and follow the default hyper-parameter setting for BlogCatalog. For other datasets, we follow the suggestion of tuning hyper-parameters from the source codes. The order is from 1 to 3, and the weights are searched according to $w_{i+1} = \beta_i w_i$ where β_i is from $\{0.01, 0.1, 1, 10, 100\}$.

FastRP [25]. We download the authors’ official source codes⁵, and follow the authors’ suggestion for hyper-parameter setting. α_1, α_2 and α_3 are set to 0, 0 and 1, respectively. We use the official tuning script to tune α_4 and the normalization strength β . The search ranges for β and α_4 are $[-1, 0]$ and $[2^{-3}, 2^6]$, respectively.

NRP [23]. We download the authors’ official source codes⁶, and follow setting in [23]: $l_1 = 20$, $l_2 = 10$, $\alpha = 0.15$, $\epsilon = 0.2$, and $\lambda = 10$.

PBG [11]. We download the authors’ official source codes⁷, and run the example script for Livejournal. For other datasets, we run the codes with default 30 iterations and report the best result.

GraphVite [12]. We adopt the reported results for YouTube, Friendster-small and Friendster in the original paper [12]. For other datasets, we run the authors’ official source codes⁸ with default setting. For Livejournal, other methods select $d = 1024$, while the official implementation of GVT⁹ only allows the selection of d up to 512.

NetSMF [7]. We download the authors’ official source codes¹⁰, and run with the default hyper-parameter setting.

ProNE [18]. We use the high-performance version of ProNE released by the LightNE GitHub Repository¹¹, and keep the

3. <https://github.com/xptree/NetMF>.

4. <https://github.com/ZW-ZHANG/RandNE>.

5. <https://github.com/GTmac/FastRP>.

6. <https://github.com/AnryYang/NRP-code>.

7. <https://github.com/facebookresearch/PyTorch-BigGraph>.

8. <https://github.com/DeepGraphLearning/graphvite>.

9. <https://graphvite.io/docs/0.2.1/api/solver.html#graphvite.solver.GraphSolver>

10. <https://github.com/xptree/NetSMF>.

11. <https://github.com/xptree/LightNE>.

TABLE 2: Statistics of datasets.

	Multi-label Vertex Classification Task				Link Prediction Task				
	BlogCatalog	YouTube	Friendster-small	Friendster	OAG	Livejournal	ClueWeb	Hyperlink2014	Hyperlink2012
$ V $	10,312	1,138,499	7,944,949	65,608,376	67,768,244	4,847,571	978,408,098	1,724,573,718	3,563,602,789
$ E $	333,983	2,990,443	447,219,610	1,806,067,142	895,368,962	68,993,773	74,744,358,622	124,141,874,032	225,840,663,232

hyper-parameters the same as those in [18]: $p = 10, \theta = 0.5, \mu = 0.2$.

LightNE [17]. We download the authors’ official source codes¹², and run experiments with the default scripts and the parameter setting according to its original paper.

SketchNE. We set parameters $b = 1, z = 8, s_1 = 100, s_2 = 1000$ for all datasets and choose T equal 2, 5 or 10, except on the very large graphs where we set $s_1 = 0$. We follow the embedding dimension d setting in [11], [12], [17], and let the other baselines follow the same setting. The eigen-decomposition rank k should be larger than d . The parameters α and q affect the experimental results to a larger extent. The bigger q , the more accurate the eigen-decomposition and thus the better the performance. We tune α and q in the experiments, choosing α in the range $[0.35, 0.5]$ with step 0.05 and q in the range $[5, 30]$ with step 5. All parameter settings for SketchNE are listed in Table. 3.

TABLE 3: Hyper-parameters for SketchNE.

Datasets	T	k	q	α	s_1	s_2	d
BlogCatalog	10	256	20	0.5	100	1000	128
YouTube	10	256	30	0.35	100	1000	128
Friendster-small	2	256	2	0.35	100	1000	128
Friendster	2	128	2	0.4	100	1000	96
OAG	10	256	30	0.45	100	1000	128
Livejournal	5	1024	10	0.35	100	1000	1024
ClueWeb	5	32	6	0.45	0	1000	32
Hyperlink2014	5	32	10	0.45	0	1000	32
Hyperlink2012	5	16	5	0.45	0	1000	16

The Setting for GNN methods. To compare SketchNE with GNN methods, we choose DGI [44] and GraphSAGE [45] as the baselines. DGI and GraphSAGE both propose to use an unsupervised loss function to train model in an unsupervised manner. Considering that GNN methods usually require a feature matrix as the input information, we follow the experimental setting in [46]. The idea is to generate non-informative features for vertices. Specifically, we generate a 128-dimensional random vector following the Xavier uniform or normal distribution for each vertice [47]. After the unsupervised training process of the GNN model converges, we obtain the embedding matrix.

The Setting for Vertex Classification. To facilitate a fair comparison, we follow the training ratio setting in [2], [7], [12], [17]. A portion of labeled vertices are sampled for training and the remaining are used for testing. We complete the task by using one-vs-rest logistic regression implemented by LIBLINEAR [48]. The prediction procedure is repeated five times and the average performance is evaluated in terms of both Micro-F1 and Macro-F1 [49].

The Setting for Link Prediction. For Livejournal, we follow the exactly same settings in Pytorch-BigGraph. For other three billion-scale graphs, we follow LightNE to set up the link prediction evaluation. We randomly excludes 0.00001% edges from the training graph for evaluation. When training SketchNE on these three graphs, the spectral propagation step is omitted due to memory cost and we set $d = 32$ except Hyperlink2012, where we use $d = 16$. We rank positive edges among randomly sampled corrupted edges to get the ranking metrics on the test set after training. We evaluate the link prediction task with four metrics—mean rank (MR), HITS@10, HITS@50, and AUC.

4.3 Experimental Results

Vertex Classification Results. We summarize the multi-label vertex classification performance in Fig. 2. In BlogCatalog, SketchNE achieves significantly better Micro-F1 and Macro-F1 than the second best method LightNE (by 3.5% on average). In YouTube [21], SketchNE show comparable performance to LightNE and GraphVite, while show significantly better results than others. In OAG [39], SketchNE achieves better performance than LightNE—the second best baseline on this data (Micro-F1 improved by 5.4% on average). In Friendster-small, and Friendster [38], SketchNE achieves the best performance among all baselines. To illustrate the effectiveness of SketchNE versus GNN methods, we tested DGI and GraphSAGE on the BlogCatalog and YouTube datasets. The non-informative features with Xavier uniform distribution or Xavier normal distribution show almost the same performance and we retain the better results between the two distributions in Fig. 2. In BlogCatalog, SketchNE achieves significantly better Micro-F1 and Macro-F1 than DGI (with 57.7% and 101.6% increase) and GraphSAGE (with 33.9% and 95.9% increase). In YouTube, SketchNE also outperforms DGI (with 73.7% and 433.6% increase) and GraphSAGE (with 89.0% and 106.2% increase) by large margins. For very large graphs, DGI and GraphSAGE cannot complete the training due to either the limitation of GPU memory size or unconvergence within a reasonable time.

Overall, SketchNE has significantly better or comparable classification results compared to other methods. Compared to RandNE, FastRP and NRP, which omit element-wise function for scalability, SketchNE shows significantly better performance. It proves that the element-wise function is crucial for learning high quality embedding and that the method in Sec. 3.1 to factorize element-wise function of low rank matrix is practical. Overall, the vertex classification results illustrate the effectiveness superiority of SketchNE.

Link Prediction Results. Table 4 lists the link prediction performance. For Livejournal, SketchNE outperforms all baselines in terms of MR, HITS@10, and HITS50. For the three billion-scale networks—Clueweb, Hyperlink2014,

12. <https://github.com/xptree/LightNE>.

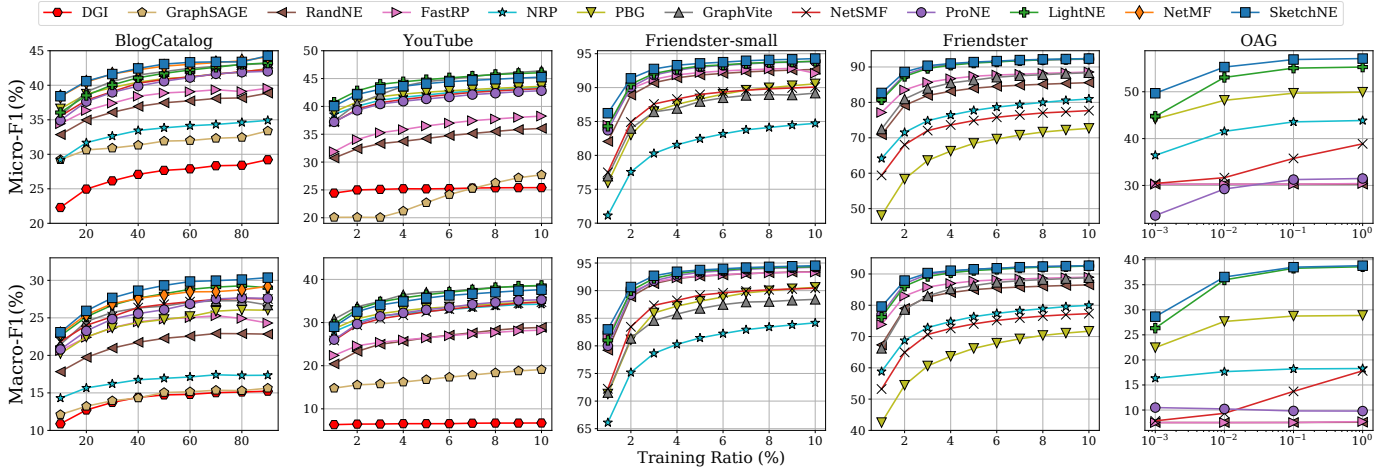


Fig. 2: Vertex classification performance (Micro-F1 and Macro-F1) w.r.t. the ratio of training data. For methods that cannot handle computation or cannot finish job in one day, the results are not available and thus not plotted in this figure.

and Hyperlink2012, we only report the results of LightNE and SketchNE, while other methods cannot finish running due to excessive memory or/and time cost. The results of LightNE are reported by choosing edge sample parameters to reach 1.5 TB memory bound. On these datasets, SketchNE produces significant outperformance over LightNE as measured by all four metrics. Take Hyperlink2012—the largest one with 3.5 billion vertices and 225 billion edges—for example, SketchNE achieves relative gains of 278%, 282%, 130%, and 24% over LightNE (the second best baseline on Livejournal) in terms of MR, HITS@10, HITS@50, and AUC.


is $2\times$ faster than LightNE and saves more than 59% memory. The results on Hyperlink2014 and Hyperlink2012 [43] further demonstrate the efficiency of SketchNE. It is worth noting that SketchNE can embed the Hyperlink2012 network with 3.5 billion vertices and 225 billion edges in 1.0 hours by using 1,321GB memory on a single machine. In conclusion, the results on the three very large-scale graphs demonstrate that SketchNE can achieve consistently and significantly better efficiency than LightNE in terms of both running time and memory cost.

TABLE 4: Link prediction comparison.

Datasets	Systems	MR ↓	HITS@10 ↑	HITS@50 ↑	AUC ↑
Livejournal	PBG	4.25	0.929	0.974	0.968
	GraphVite	3.06	0.962	0.982	0.973
	NetSMF	3.09	0.954	0.986	0.935
	RandNE	5.19	0.912	0.966	0.957
	FastRP	4.51	0.928	0.973	0.965
	NRP	2.98	0.949	0.993	0.934
	ProNE	3.31	0.950	0.982	0.932
	LightNE	2.13	0.977	0.993	0.945
	SketchNE	2.10	0.977	0.994	0.945
	ClueWeb	LightNE	105.9	0.753	0.803
SketchNE		32.0	0.771	0.869	0.968
Hyperlink2014	LightNE	129.7	0.5	0.628	0.874
	SketchNE	110.3	0.593	0.693	0.890
Hyperlink2012	LightNE	257.7	0.189	0.348	0.751
	SketchNE	68.1	0.722	0.802	0.933

Time and Memory Efficiency. We report the running time and memory cost of SketchNE and other eight baselines on all nine datasets in Table 5. Time-wise, on small datasets with millions of edges, the running time of SketchNE is relatively comparable to other baselines. However, on networks of billions of edges, e.g., Friendster, ClueWeb, Hyperlink2014 and Hyperlink2012, it takes SketchNE the least time to embed them. Memory-wise, we can observe that SketchNE demands less memory than all baselines on all datasets except using slightly more memory than NetSMF and ProNE on the small Livejournal data, empowering it to go for the largest networks considered and beyond. For example, the running time of SketchNE on ClueWeb [42] is 37.7 minutes and the peak memory cost is 612GB, which

4.4 Ablation and Case Studies



(a) The computed eigenvalues w.r.t. α . (b) NetMF vs. NetMF (w/ freigs) α . **Fig. 3:** The validation of the effectiveness of freigs.

Efficiency and Effectiveness of Each Step of SketchNE. First, we focus on the fast randomized eigen-decomposition (*freigs*, Alg. 6 Line 1). Halko et al. [29] has shown that it is challenging to perform eigen-decomposition on $D^{-1/2}AD^{-1/2}$. We choose YouTube as an example, whose related eigenvalues are shown in Fig. 3 (a). The eigenvalues of $D^{-1/2}AD^{-1/2}$ with power iteration $q = 20$ and $q = 100$ are far from correct, as the correct largest eigenvalue should be 1. The *eigsh* [28] is not able to complete this job in three days. The results illustrate the requirement to use modified Laplacian matrix for fast randomized eigen-decomposition. When we choose $\alpha = 0.4$, the eigenvalues of fast randomized eigen-decomposition ($q = 20$) on $D^{-0.4}AD^{-0.4}$ is indistinguishable from the eigenvalues computed by

TABLE 5: Efficiency comparison.

Metric	Datasets	RandNE	FastRP	NRP	PBG	GraphVite	NetSMF	ProNE	LightNE	SketchNE
Time	BlogCatalog	0.5 s	1.0 s	3.0 s	174.0 s	4.0 s	11.3 m	79.0 s	152.0 s	2.0 s
	YouTube	12.0 s	17.0 s	173.0 s	12.5 m	44.0 s	3.7 h	65.0 s	96.0 s	40.0 s
	Friendster-small	11.8 m	40.0 m	3.5 h	22.7 m	2.8 h	52 m	5.3 m	7.5 m	5.2 m
	Friendster	57.8 m	3.5 h	16.1 h	5.3 h	20.3 h	16.5 h	19.5 m	37.6 m	16.0 m
	OAG	33.7 m	3.5 h	11.4 h	20 h	1+day	22.4 h	22.6 m	1.5 h	1.1 h
	Livejournal	9.0 m	18.0 m	4.3 h	7.3 h	29.0 m	2.1 h	12.8 m	16.0 m	12.5 m
	ClueWeb	×	×	×	1+day	1+day	×	×	1.3 h	37.7 m
	Hyperlink2014	×	×	×	1+day	1+day	×	×	1.8 h ¹	0.98 h
	Hyperlink2012	×	×	×	1+day	1+day	×	×	5.6 h ¹	1.0 h
	Mem (GB)	BlogCatalog	0.2	0.3	0.6	*	*	135	18	273
YouTube		6.9	12.5	9.7	*	*	854	28	83	27
Friendster-small		125	125	105	*	*	85	84	541	56
Friendster		548	583	400	*	*	1144	326	559	236
OAG		429	746	473	*	*	1500	403	1391	283
Livejournal		224	417	282	*	*	140	131	532	147
ClueWeb		×	×	×	*	*	×	×	1493	612
Hyperlink2014		×	×	×	*	*	×	×	1500 ¹	1076
Hyperlink2012		×	×	×	*	*	×	×	1500 ¹	1321

¹ LightNE requires sufficient samples that cost more than 1500GB mem (and more time), so it has to stop before it reaches the mem limit. “*” indicates that we do not compare the memory cost of the CPU-GPU hybrid system (GraphVite) or distributed memory system (PBG). “×” indicates that the corresponding algorithm is unable to handle the computation due to excessive space and memory consumption.

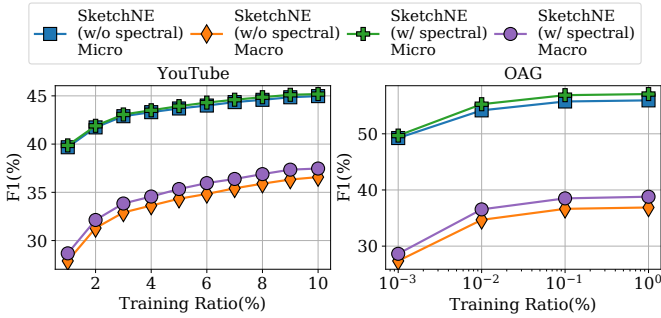


Fig. 4: The embedding performance comparison.

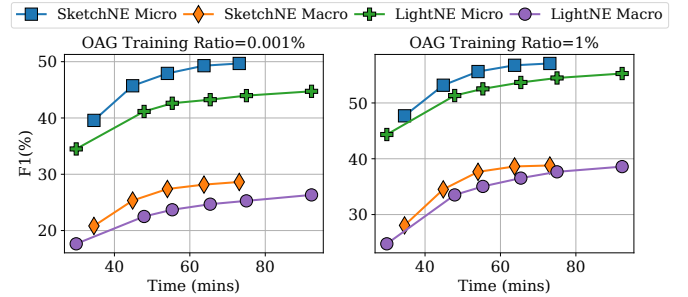


Fig. 5: The trade-offs between efficiency and performance.

eigsh when $k < 140$. The running time of fast randomized eigen-decomposition is 20 seconds while eigsh costs 31 minutes, which proves the accuracy and efficiency of fast randomized eigen-decomposition. Then, we replace the eigen-decomposition (eigsh) of NetMF with Alg. 5 and evaluate the Micro-F1 result of BlogCatalog between NetMF and NetMF (w/ freigs) in Fig. 3 (b). The results show the effectiveness of freigs. Next, we focus on the effects of sparse-sign randomized single-pass SVD and spectral propagation. SketchNE (w/o spectral) and SketchNE (w/ spectral) represent the result of initial and enhanced embeddings, respectively. The vertex classification result is shown in Fig. 4. The performance of model (w/o spectral) has been satisfactory, which proves the effectiveness of the sparse-sign randomized single-pass SVD. Combining with spectral propagation, SketchNE (w/ spectral) shows better results, demonstrating the effect of spectral propagation.

The effects of parameter q . Here we need to pay attention to q , which determines the accuracy of the fast randomized eigen-decomposition. Thus, we make a trade-off between the quality of the learned embeddings and the overall running time. The peak memory cost is constant when we fix the oversampling parameters, col-

umn density z , eigen-decomposition rank k , and embedding dimension d . By fixing other parameters for OAG, we enumerate q from $\{10, 15, 20, 25, 30\}$, the efficiency-effectiveness trade-off of OAG is shown as Fig. 5. We also add LightNE results with different edge samples M from $\{1Tm, 5Tm, 7Tm, 10Tm, 13Tm, 17Tm\}$. The peak memory of SketchNE is still 283GB while that of LightNE is 553 GB, 682 GB, 776 GB, 936 GB, 1118 GB and **1391 GB**, respectively. Fig. 5 shows that SketchNE can learn more effective embeddings than LightNE when the running time is constant with less memory cost. The experiment proves that users can adjust SketchNE flexibly according to time/memory budgets and performance requirements.

The effects of parameter α . We also analyze the influence of parameter α , which balances the accuracy of fast randomized eigen-decomposition on modified Laplacian matrix and the approximation error of Eq. (11). We select here YouTube and OAG as example datasets for the ablation study. We vary the parameter α from $\{0.25, 0.3, 0.35, 0.4, 0.45, 0.5\}$ and fix other parameters. The change of SketchNE’s performance as the α varies can be shown in Fig. 6. From it we can see that when setting $\alpha = 0.5$, the eigen-decomposition on the Laplacian matrix has a bad accuracy and therefore causes a loss of performance in Fig. 6. Fig. 6 also shows




Fig. 6: The performance change when varying parameter α .

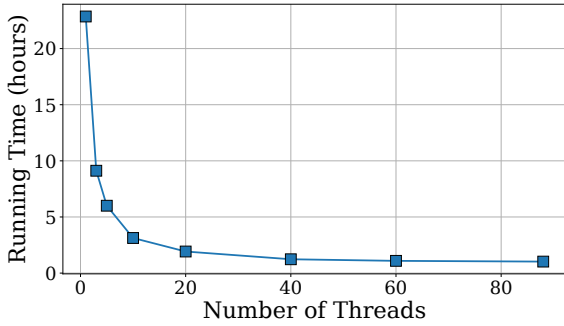


Fig. 7: Runtime of SketchNE v.s. the number of threads.

that SketchNE can learn almost the best embeddings when setting α in the range [0.35, 0.45].

The Sketching Method. We download the official source codes¹³ of Polynomial Tensor Sketch [31], and set $k = 10, r = 11, m = 12$, which makes dimension $m \times r + 1$ slightly bigger than SketchNE’s $d = 128$ for BlogCatalog. We replace the sparse-sign randomized single-pass SVD in SketchNE with the polynomial tensor sketch algorithm, which is an alternative solution to factorize $f^\circ(LR)$. Experiments on BlogCatalog show that SketchNE with sparse-sign randomized single-pass SVD performs much better than SketchNE with polynomial tensor sketch, achieving a 44%/67.8% relative improvement on Micro-F1/Macro-F1 with 10% of training data.

The Number of Threads. In this work, we use a single-machine shared memory implementation with multi-threading acceleration. We set the number of threads to be 1, 3, 5, 10, 20, 40, 60, 88, and report the corresponding running time of SketchNE in Fig. 7. SketchNE takes 22.8 hours to embed the Hyperlink2012 network with 1 threads and 1.9 hours with 40 threads, achieving a 12 \times speedup ratio (with ideal being 40 \times). This relatively good sub-linear speedup supports SketchNE to scale up to networks with hundreds of billions of edges.

5 RELATED WORK

This section reviews the related work of network embedding and randomized matrix factorization.

Network Embedding. Network embedding has been comprehensively studied over the past decade. Recent work

about network embedding can be divided into three categories. The first category is based on skip-gram methods, inspired by word2vec [10], including DeepWalk [2], Line [4], node2vec [20]. These methods rely on stochastic gradient descent to optimize a logistic loss. The second category is based on deep learning methods, including GCN [50], GAT [51], GIN [52], ie-HGCN [46], GCC [53], GraphCL [54], DGI [44] and GraphSAGE [45]. Here we only focus on graphs with no additional information and pre-train general network embeddings in an unsupervised manner. GraphSAGE [45] and DGI [44] propose a loss function which can be used to train GNN models in an unsupervised manner. GCC [53] propose a framework to capture the universal network topological properties across multiple networks. GraphCL [54] improves the performance of graph contrastive learning by data augmentations. ie-HGCN [46] presents an interpretable and efficient heterogeneous graph convolutional network. In contrast, these methods pre-train in a supervised manner with vertex attributes. The third category method is based on matrix factorization, using SVD to generate the best low-rank approximation [55]. GraRep [56], HOPE [57], NetMF [13], NetSMF [7], ProNE [18], NRP [23] and LightNE [17] are methods in this category, while PANE [58] is designed for networks with attributes. There are several high performance embedding systems for large graphs have been developed. GraphVite [12], a CPU-GPU hybrid network embedding system, is developed based on DeepWalk [2] and LINE [4]. In GraphVite, CPU is used to perform graph operation, and GPU is used to compute linear algebra. Nevertheless, the GPU memory is a disadvantage when processing billion scale networks, limiting widespread use. Based on DeepWalk [2] and LINE [4], PyTorch-BigGraph [11] has been proposed for distributed memory machines. It achieves load balancing by graph partition and synchronization through a shared parameter server. In this work, we propose SketchNE, which leverages the merit of NetMF and addresses its limitation in speed and memory overhead. Then spectral propagation is applied to improve the quality of final embedding.

Randomized Matrix Factorization. As the amount of data continues to increase, the popularity of randomized matrix computations has grown significantly. Randomized SVD can be an alternative to conventional SVD methods, because it involves the same or fewer floating-point operations and is more efficient for truly large high-dimensional data, by exploiting modern computing architectures [29]. As an application, frPCA [32] is developed for large-scale sparse data with better performance than conventional PCA algorithms. Randomized SVD has shown superiority in recent network embedding methods such as NetSMF, ProNE, and LightNE. Over the past few years, several single-pass SVD algorithms [30], [59] based on randomized matrix sketch have been introduced for streaming data scenarios. The efficiency and memory limitation of NetMF will be solved by randomized matrix factorization.

6 CONCLUSION

In this work, we propose SketchNE, a fast, memory-efficient, and scalable network embedding method. We formulate

13. <https://github.com/insuhan/polytensorsketch>.

the computation goal of NetMF as factorizing an element-wise function of low-rank matrix and then analyze its computational challenges. SketchNE resolves these challenges by leveraging various randomized linear sketch techniques, including but not limited to sparse-sign matrix, single-pass SVD, and fast eigen-decomposition. We use sparse-sign random projection matrix to solve the matrix multiplication challenge between $f^\circ(\mathbf{LR})$ and random projection matrix, and generate a sketch that can capture the dominant information of $f^\circ(\mathbf{LR})$. Then, we solve the challenge in constructing reduced matrix by single-pass SVD. Secondly, we propose a fast randomized eigen-decomposition algorithm for modified Laplacian matrix. To enhance the performance of embedding, spectral propagation is adopted and a high-performance parallel graph processing stack GBBS is used to achieve memory-efficiency. The main computation steps of SketchNE are highly parallelizable, which is thus well supported by the MKL library and OpenMP. With the help of these techniques, SketchNE achieves the best performance on vertex classification and link prediction among state-of-the-art methods across diverse datasets. Notably, SketchNE can learn high-quality embeddings for a network with 3.5 billion vertices and 225 billion edges in 1.0 hour by using 1,321GB memory on a single machine, and the learned embeddings offer a 282% relative HITS@10 improvement over LightNE on the link prediction task.

In the future, we plan to extend the method to handle dynamic and heterogeneous networks. Since the sparse-sign randomized single-pass SVD is proposed for solving the problem $f^\circ(\mathbf{LR})$, it is foreseeable that Alg. 4 may have advantages in approximating the basic components in deep learning network.

REFERENCES

- [1] W. L. Hamilton, "Graph Representation Learning," *Synthesis Lectures on Artificial Intelligence and Machine Learning*, vol. 14, no. 3, pp. 1–159, 2020.
- [2] B. Perozzi, R. Al-Rfou, and S. Skiena, "Deepwalk: Online Learning of Social Representations," in *KDD '14*, 2014, pp. 701–710.
- [3] J. Wang, P. Huang, H. Zhao, Z. Zhang, B. Zhao, and D. L. Lee, "Billion-Scale Commodity Embedding for E-commerce Recommendation in Alibaba," in *KDD '18*, 2018, pp. 839–848.
- [4] J. Tang, M. Qu, M. Wang, M. Zhang, J. Yan, and Q. Mei, "LINE: Large-Scale Information Network Embedding," in *WWW '15*, 2015.
- [5] R. Ramanath, H. Inan, G. Polatkan, B. Hu, Q. Guo, C. Ozcaglar, X. Wu, K. Kenthapadi, and S. C. Geyik, "Towards Deep and Representation Learning for Talent Search at LinkedIn," in *CIKM '18*, 2018.
- [6] Y. Dong, N. V. Chawla, and A. Swami, "metapath2vec: Scalable representation learning for heterogeneous networks," in *KDD '17*. ACM, 2017, pp. 135–144.
- [7] J. Qiu, Y. Dong, H. Ma, J. Li, C. Wang, K. Wang, and J. Tang, "NetSMF: Large-Scale Network Embedding as Sparse Matrix Factorization," in *WWW '19*, 2019, pp. 1509–1520.
- [8] M. Academic, "Multi-Sense Network Representation Learning in Microsoft Academic Graph," <https://www.microsoft.com/en-us/research/project/academic/articles/multi-sense-network-representation-learning-in-microsoft-academic-graph/>, 2020, [Online; accessed Nov. 21, 2022].
- [9] R. Ying, R. He, K. Chen, P. Eksombatchai, W. L. Hamilton, and J. Leskovec, "Graph Convolutional Neural Networks for Web-Scale Recommender Systems," in *KDD '18*, 2018, p. 974–983.
- [10] T. Mikolov, K. Chen, G. Corrado, and J. Dean, "Efficient Estimation of Word Representations in Vector Space," in *ICLR '13*, 2013.
- [11] A. Lerer, L. Wu, J. Shen, T. Lacroix, L. Wehrstedt, A. Bose, and A. Peysakhovich, "PyTorch-BigGraph: A Large-Scale Graph Embedding System," *SysML '19*, 2019.
- [12] Z. Zhu, S. Xu, J. Tang, and M. Qu, "GraphVite: A High-Performance CPU-GPU Hybrid System for Node Embedding," in *WWW '19*, 2019, pp. 2494–2504.
- [13] J. Qiu, Y. Dong, H. Ma, J. Li, K. Wang, and J. Tang, "Network Embedding as Matrix Factorization: Unifying Deepwalk, LINE, PTE, and node2vec," in *WSDM '18*, 2018, pp. 459–467.
- [14] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention Is All You Need," in *NeurIPS '17*, 2017, pp. 5998–6008.
- [15] X. Glorot, A. Bordes, and Y. Bengio, "Deep Sparse Rectifier Neural Networks," in *AISTATS '11*, 2011, pp. 315–323.
- [16] T. Hofmann, B. Schölkopf, and A. J. Smola, "Kernel Methods in Machine Learning," *The annals of statistics*, vol. 36, no. 3, pp. 1171–1220, 2008.
- [17] J. Qiu, L. Dhulipala, J. Tang, R. Peng, and C. Wang, "LightNE: A Lightweight Graph Processing System for Network Embedding," in *SIGMOD '21*, 2021, pp. 2281–2289.
- [18] J. Zhang, Y. Dong, Y. Wang, J. Tang, and M. Ding, "ProNE: Fast and Scalable Network Representation Learning," in *IJCAI '19*, 2019, pp. 4278–4284.
- [19] L. Dhulipala, G. E. Blelloch, and J. Shun, "Theoretically Efficient Parallel Graph Algorithms Can Be Fast and Scalable," in *SPAA '18*, 2018, pp. 393–404.
- [20] A. Grover and J. Leskovec, "node2vec: Scalable Feature Learning for Networks," in *KDD '16*, 2016, pp. 855–864.
- [21] L. Tang and H. Liu, "Scalable Learning of Collective Behavior Based on Sparse Social Dimensions," in *CIKM '09*, 2009, pp. 1107–1116.
- [22] V. Ivashkin and P. Chebotarev, "Do Logarithmic Proximity Measures Outperform Plain Ones in Graph Clustering?" in *International Conference on Network Analysis*. Springer, 2016, pp. 87–105.
- [23] R. Yang, J. Shi, X. Xiao, Y. Yang, and S. S. Bhowmick, "Homogeneous Network Embedding for Massive Graphs via Reweighted Personalized Pagerank," *VLDB '20*, p. 670–683, 2020.
- [24] Z. Zhang, P. Cui, H. Li, X. Wang, and W. Zhu, "Billion-Scale Network Embedding With Iterative Random Projection," in *ICDM '18*, 2018.
- [25] H. Chen, S. F. Sultan, Y. Tian, M. Chen, and S. Skiena, "Fast and Accurate Network Embeddings via Very Sparse Random Projection," in *CIKM '19*, 2019, pp. 399–408.
- [26] D. Cohen-Steiner, W. Kong, C. Sohler, and G. Valiant, "Approximating the Spectrum of a Graph," in *KDD '18*, 2018, pp. 1263–1271.
- [27] J. Liu, C. Wang, M. Danilevsky, and J. Han, "Large-Scale Spectral Clustering on Graphs," in *IJCAI '13*, 2013.
- [28] R. B. Lehoucq, D. C. Sorensen, and C. Yang, *Arpack Users' Guide: Solution of Large-Scale Eigenvalue Problems With Implicitly Restarted Arnoldi Methods*. SIAM, 1998.
- [29] N. Halko, P.-G. Martinsson, and J. A. Tropp, "Finding Structure With Randomness: Probabilistic Algorithms for Constructing Approximate Matrix Decompositions," *SIAM Review*, vol. 53, no. 2, pp. 217–288, 2011.
- [30] J. A. Tropp, A. Yurtsever, M. Udell, and V. Cevher, "Streaming Low-Rank Matrix Approximation With an Application to Scientific Simulation," *SIAM Journal on Scientific Computing*, vol. 41, no. 4, pp. A2430–A2463, 2019.
- [31] I. Han, H. Avron, and J. Shin, "Polynomial Tensor Sketch for Element-Wise Function of Low-Rank Matrix," in *ICML '20*, 2020, pp. 3984–3993.
- [32] X. Feng, Y. Xie, M. Song, W. Yu, and J. Tang, "Fast Randomized PCA for Sparse Data," in *ACML '18*, 2018, pp. 710–725.
- [33] J. Shun and G. E. Blelloch, "Ligra: A Lightweight Graph Processing Framework for Shared Memory," in *ACM Sigplan Notices*, vol. 48, no. 8, 2013, pp. 135–146.
- [34] J. Kepner and J. Gilbert, *Graph Algorithms in the Language of Linear Algebra*. SIAM, 2011.
- [35] J. Shun, L. Dhulipala, and G. E. Blelloch, "Smaller and Faster: Parallel Processing of Compressed Graphs With Ligra+," in *DCC '15*, 2015, pp. 403–412.
- [36] L. Dagum and R. Menon, "OpenMP: An Industry-Standard API for Shared-Memory Programming," vol. 5, no. 1, p. 46–55, 1998.
- [37] L. Tang and H. Liu, "Relational Learning via Latent Social Dimensions," in *KDD '09*, 2009, pp. 817–826.
- [38] J. Yang and J. Leskovec, "Defining and Evaluating Network Communities Based on Ground-Truth," *KAIS '15*, vol. 42, no. 1, pp. 181–213, 2015.

- [39] A. Sinha, Z. Shen, Y. Song, H. Ma, D. Eide, B.-J. P. Hsu, and K. Wang, "An Overview of Microsoft Academic Service (MAS) And Applications," in *WWW '15 Companion*, 2015, p. 243–246.
- [40] J. Tang, J. Zhang, L. Yao, J. Li, L. Zhang, and Z. Su, "Arnetminer: Extraction and Mining of Academic Social Networks," in *KDD '08*, 2008, p. 990–998.
- [41] J. Leskovec, K. J. Lang, A. Dasgupta, and M. W. Mahoney, "Community Structure in Large Networks: Natural Cluster Sizes and the Absence of Large Well-Defined Clusters," *Internet Mathematics*, vol. 6, no. 1, pp. 29–123, 2009.
- [42] P. Boldi and S. Vigna, "The WebGraph Framework I: Compression Techniques," in *WWW '04*, 2004, pp. 595–601.
- [43] R. Meusel, S. Vigna, O. Lehmborg, and C. Bizer, "The Graph Structure in the Web – Analyzed on Different Aggregation Levels," *The Journal of Web Science*, vol. 1, 2015.
- [44] P. Velickovic, W. Fedus, W. L. Hamilton, P. Liò, Y. Bengio, and R. D. Hjelm, "Deep graph infomax." *ICLR'19*, 2019.
- [45] W. Hamilton, Z. Ying, and J. Leskovec, "Inductive representation learning on large graphs," *NeurIPS '17*, vol. 30, 2017.
- [46] Y. Yang, Z. Guan, J. Li, W. Zhao, J. Cui, and Q. Wang, "Interpretable and efficient heterogeneous graph convolutional network," *IEEE Transactions on Knowledge and Data Engineering*, 2021.
- [47] X. Glorot and Y. Bengio, "Understanding the difficulty of training deep feedforward neural networks," in *Proceedings of the thirteenth international conference on artificial intelligence and statistics. JMLR Workshop and Conference Proceedings*, 2010, pp. 249–256.
- [48] R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin, "LIBLINEAR: A Library for Large Linear Classification," *JMLR '08*, vol. 9, pp. 1871–1874, 2008.
- [49] G. Tsoumakas, I. Katakis, and I. Vlahavas, "Mining Multi-Label Data," in *Data Mining and Knowledge Discovery Handbook*, 2009, pp. 667–685.
- [50] T. N. Kipf and M. Welling, "Semi-supervised Classification With Graph Convolutional Networks," in *ICLR '17*, 2017.
- [51] P. Velickovic, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio, "Graph Attention Networks," in *ICLR '18*, 2018.
- [52] K. Xu, W. Hu, J. Leskovec, and S. Jegelka, "How Powerful Are Graph Neural Networks?" in *ICLR '19*, 2019.
- [53] J. Qiu, Q. Chen, Y. Dong, J. Zhang, H. Yang, M. Ding, K. Wang, and J. Tang, "Gcc: Graph contrastive coding for graph neural network pre-training," in *KDD'20*, 2020, pp. 1150–1160.
- [54] Y. You, T. Chen, Y. Sui, T. Chen, Z. Wang, and Y. Shen, "Graph contrastive learning with augmentations," *NeurIPS '20*, vol. 33, pp. 5812–5823, 2020.
- [55] C. Eckart and G. Young, "The Approximation of One Matrix by Another of Lower Rank," *Psychometrika*, vol. 1, no. 3, pp. 211–218, 1936.
- [56] S. Cao, W. Lu, and Q. Xu, "GraRep: Learning Graph Representations With Global Structural Information," in *CIKM '15*, 2015, pp. 891–900.
- [57] M. Ou, P. Cui, J. Pei, Z. Zhang, and W. Zhu, "Asymmetric Transitivity Preserving Graph Embedding," in *KDD '16*, 2016, p. 1105–1114.
- [58] R. Yang, J. Shi, X. Xiao, Y. Yang, J. Liu, and S. S. Bhowmick, "Scaling attributed network embedding to massive graphs," *Proceedings of the VLDB Endowment*, vol. 14, no. 1, pp. 37–49, 2020.
- [59] W. Yu, Y. Gu, J. Li, S. Liu, and Y. Li, "Single-Pass PCA of Large High-Dimensional Data," in *IJCAI '17*, 2017, p. 3350–3356.