

Delving Into Deep Walkers: A Convergence Analysis of Random-Walk-Based Vertex Embeddings

Dominik Kloefer, Angelica I. Aviles-Rivero and Daniel Heydecker

Abstract—Graph vertex embeddings based on random walks have become increasingly influential in recent years, showing good performance in several tasks as they efficiently transform a graph into a more computationally digestible format while preserving relevant information. However, the theoretical properties of such algorithms, in particular the influence of hyperparameters and of the graph structure on their convergence behaviour, have so far not been well-understood. In this work, we provide a theoretical analysis for random-walks based embeddings techniques. Firstly, we prove that, under some weak assumptions, vertex embeddings derived from random walks do indeed converge both in the single limit of the number of random walks $N \rightarrow \infty$ and in the double limit of both N and the length of each random walk $L \rightarrow \infty$. Secondly, we derive concentration bounds quantifying the converge rate of the corpora for the single and double limits. Thirdly, we use these results to derive a heuristic for choosing the hyperparameters N and L . We validate and illustrate the practical importance of our findings with a range of numerical and visual experiments on several graphs drawn from real-world applications.

Index Terms—Machine Learning, Graph Embeddings, Vertex Embedding, Convergence, Random Walk, Markov Chain, DeepWalk.

1 INTRODUCTION

GRAPHS naturally represent data arising in several real-world scenarios; important examples include social networks, medical records, protein networks and the web. In many tasks, using graph embeddings that meaningfully encode relevant information about the graph structure has enabled outstanding performance on downstream tasks, for example in node classification [1], link prediction [2], anomaly detection [3], and node clustering [4], [5].

This has motivated the development of several techniques which generate embeddings by digesting the graph in different ways. One set of techniques is based on matrix factorisation, e.g. [6], [7], [8], [9], [10], in which the connections between nodes (graph properties) are represented in matrix form to obtain the embeddings. Other bodies of research have explored approaches including generative models [11], [12], [13] and hybrid (i.e. a combination of) techniques [14], [15], [16], to name a few.

Of particular interest for this paper are random-walk-based techniques, e.g. [18], [20], [21], [22], who have become popular due to their good performance. The core idea is to create a *context* for each vertex by generating sequences of vertices using random walks on the graph before generating vertex embeddings from these contexts. Motivated by the success of neural language models, starting with the seminal DeepWalk [18] and followed by others including node2vec [20] and DNDR [22], a variety of algorithms view individual random walks as “sentences” consisting

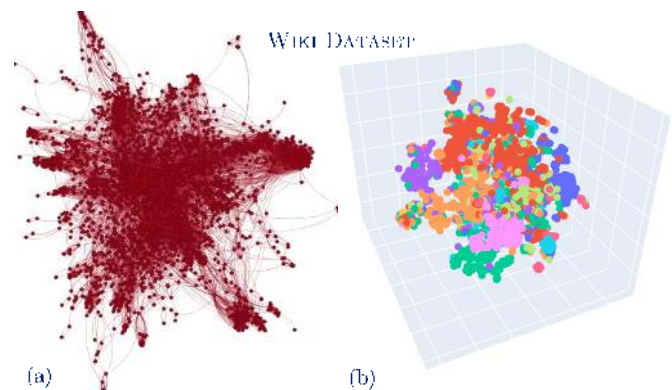


Fig. 1. Input-output of a random walk based vertex embeddings technique from one of our experiments. (a) Input graph from the Wiki dataset [17]. (b) Learned latent vertex representation generated by DeepWalk [18]. Colors indicate ground-truth classes in the dataset, and the high-dimensional embeddings have been visualised using t-SNE [19].

of individual “words” (individual vertices); a collection of random walks then corresponds to a “text corpus”. A corpus of random walks derived from a graph during a first stage can then be analysed in a second stage using language models such as SkipGram [23], which creates word (vertex) embeddings that allow one to predict a word’s context from its embedding. This process is illustrated in Fig. 1 using one of our experiments, where the graph input and the latent vertex representation are displayed.

Contributions. Despite their wide-spread usage, not much effort has been spent on developing a firm theoretical understanding of these algorithms. In particular, the study

- D. Kloefer is with the Department of Engineering, University of Oxford, Oxford OX1 3PJ.
E-mail: dominik.kloefer@eng.ox.ac.uk
- AI Aviles-Rivero and D. Heydecker are with the Department of Applied Mathematics and Theoretical Physics, University of Cambridge, Cambridge CB3 0WA.
E-mail: {ai323, dh489}@cam.ac.uk

of the convergence of random-walk based algorithms is relatively new and incomplete. While there are several works that slightly vary the graph sampling procedure in the first stage, e.g. [20] and [24], in this paper we mainly consider algorithms that share the sampling strategy of DeepWalk [18] due to this strategy’s practical advantages (requiring, for example, fewer hyperparameters than node2vec [20]) and because it continues to be the backbone for many new techniques and applications, e.g. [25], [26], [27], [28]. Our contributions are:

- After formally describing DeepWalk-related techniques (Section 3), we provide a theoretical analysis for random-walk based embedding algorithms. (Section 4)
 - In Section 4.1, we formalise the random-walk sampling strategy used by DeepWalk [18] in its first stage and derive the asymptotic limit of the corpora generated. We show that as the number of random walks N goes to ∞ this convergence is almost sure both in the single limit $N \rightarrow \infty$ and the double limit $N, L \rightarrow \infty$ where L is the length of each random walk. (Theorem 1, Corollary 1.1)
 - In Section 4.2, we derive concentration bounds that quantify the convergence rate of the corpora in both these cases. (Theorem 2, Theorem 6)
 - We apply the results of this analysis to derive a heuristic that gives optimal values for the hyperparameters N and L for a fixed computational cost. (Corollary 6.1)
- We validate our theoretical results through a set of experiments using medium-sized and large real-world graphs consisting of thousands to millions vertices / edges. (Section 5)

2 RELATED WORK

This paper is interested in random-walk-based techniques for vertex embeddings, which have demonstrated outstanding performance. We first briefly review existing such techniques.

Random-Walk Based Vertex Embeddings. The seminal work of Perozzi et al. [18] introduced the DeepWalk technique, which learns embeddings in two stages. In the first stage, a given graph is sampled using random walks, followed by an application of the SkipGram model to then generate vertex embeddings that allow one to predict the vertices that appear within a window around the vertex in the random walks. The promising performance of the DeepWalk motivated the developments of subsequent techniques that follow the principles of the DeepWalk.

Grover et al. [20] introduce an additional parameter to control the random walks by allowing the user to interpolate between breadth-first and depth-first graph searches. Another modification of the first stage called Walklets was introduced in [24], where authors skipped over steps in each random walk. The HARP model introduced in [29] improves the initialisation of the SkipGram stage. Slightly improved performance of these models compared to DeepWalk however comes at the cost of increased computational costs.

Another set of techniques modified DeepWalk by introducing different sampling techniques for the random walks (e.g. [21], [30]), and enforcing preservation of some graph properties (e.g. [31], [32], [33]). Other variants of the such family of techniques are the works of [9], [22], [34], [35]. There exist also techniques that aim to create similar network embeddings without (explicitly) sampling any random walks, for example [36].

While DeepWalk and its variants have been widely explored in the literature from a practical point of view, the analysis of the theoretical properties of such techniques is still relatively new. There are only few works that discussed the theoretical properties of this family of techniques. In the following, we discuss them and note the differences with the present work.

Theoretical Analyses of DeepWalk & Comparison to Our Work. The first theoretical analysis of random-walk based vertex embedding techniques is due to Qiu et al. [26], who demonstrated a relationship between DeepWalk and a particular matrix factorisation problem. In the notation of Section 3, that work identifies the limiting behaviour as the length of the walk $L \rightarrow \infty$, but does not prove a quantitative rate, and keeps the number of trials N fixed rather than identifying the convergence as $N \rightarrow \infty$ or the convergence of the joint limit as done here. The behaviour of DeepWalk in the limit of many shorter (instead of fewer long) random walks as investigated in our work arguably has greater practical relevance as it allows for much easier parallelisation of the algorithm. Additionally, the proof in [26] only applies to undirected, non-bipartite graphs whereas our Theorem 1 and Corollaries 1.1 and 1.2 apply to directed and undirected graphs, both periodic and aperiodic.

More recently, another work [37] proved a concentration bound in terms of L for arbitrary N for the convergence of a corpus of random walks generated from an aperiodic graph. In contrast, we derive concentration bounds in terms of N for arbitrary L (Theorem 2) and in terms of N and L for the joint limit (Theorem 6) for both periodic and aperiodic graphs. Furthermore, our bounds do not depend on the size of the graph, a factor that while of constant order can be very significant in real-world graphs with vertices and edges potentially numbering in the millions (cf. Table 1). In this paper, we also consider the assumptions necessary for convergence of the co-occurrence matrix to imply convergence of the extracted vertex representations (Theorem 7), a question that [37] does not address.

Most recently, [38] considers random graphs, in either the sparse or dense regimes, in the asymptotic where the size $|V|$ and the number of sampled paths $N \rightarrow \infty$ on a fixed time interval. While also investigating the theoretical properties of random-walk based graph embedding algorithms, the setting that they consider is different to the one considered in this paper or, in fact, in [37].

Moreover and in contrast to the works above, we illustrate the practical relevance of our analysis in Corollary 6.1 by combining our results to derive a rule for the optimal choice of the hyperparameters N and L . Qiu et al. [37] also derive a heuristic for choosing L , but they assume implicitly that $N = 1$, which as our experiments in Section 5 show yields worse task-performance than our heuristic.

3 FORMAL DESCRIPTION OF THE DEEPWALK ALGORITHM

The DeepWalk algorithm consists of two phases: in the first phase it generates a corpus of vertex pairs, and in the second phase it executes the SkipGram algorithm (with negative sampling) on this corpus to generate vertex embeddings.

3.1 Generation of Corpus

Informally, in the first phase the corpus \mathcal{D} is populated with all pairs of vertices that are within a given window size T of each other in random walks generated on the graph. For each vertex, the vertices at most T removed in the random walk are viewed as that vertex's "context" vertices, and each (vertex, context vertex) pair is added to the corpus.

We take a graph $G = (V, E)$, equipped with a weight function $w : E \rightarrow [0, \infty)$. Consider a random walk $(V_t)_{t \geq 0}$ on G , with a given starting distribution f_{V_0} and transition matrix P with elements $P_{ij} = w_{ij} / \sum_j w_{ij}$, where w_{ij} is the weight of edge (v_i, v_j) . We write $\text{MARKOV}(f_{V_0}, P, L)$ for the joint law of the first L steps (V_0, \dots, V_{L-1}) and, if $f_{V_0} = \delta_{v_0}$ is a one-hot distribution, write $\text{MARKOV}(v_0, P, L)$ instead. The stationary distribution of this random walk is denoted as $\pi \in \mathbb{R}^{|V|}$, which is given explicitly by $\pi(v_i) \propto \sum_j w_{ij}$.

We define a *corpus* \mathcal{D} on graph $G = (V, E)$ to be a tuple (D, m) consisting of the set of pairs $D = V \times V$ and the *multiplicity function* $m : D \rightarrow \mathbb{N}$. Its *cardinality* is the total multiplicity $|\mathcal{D}| = \sum_{d \in D} m(d)$. Since $|\mathcal{D}|$ is finite, we may identify m with the vector of its values, and write $m \in \mathbb{R}_{\geq 0}^{|V| \times |V|}$, and we denote its elements as $m(i, j) = m((v_i, v_j))$ for all $v_i, v_j \in V$, so for all $(v_i, v_j) \in D$. This matrix m is also called the *co-occurrence matrix*, as it encodes the number of co-occurrences of two vertices (the number of times the appear within T random walk steps of each other) in the random walks performed. Note that once we have specified a graph, different corpora only differ in their multiplicity functions.

In the first phase of DeepWalk, an initially empty corpus is populated by repeatedly updating the multiplicity function. Formally, the algorithm with objective to populate the corpus \mathcal{D} is given in Algorithm 1, taking as parameters the number of walks N , the walk length L and window size T .

In the second loop in line 5, we follow [26] in having j take on $L - T$ different values to avoid edge effects (when $j + r > L - 1$). This simplifies the derivation in Section 4.1 while only removing constant order terms (depending only on the fixed length T) from the equations.

3.2 Generation of Vertex Representation

Having obtained a corpus in the first stage, in the second phase the vertex representations are calculated from the corpus. Traditionally using either the SkipGram algorithm [23] or SkipGram with Negative Sampling [39] (though in principle other methods can be used as well), the vertex representations $\{z_i\}_{i: v_i \in V}$ are obtained by maximising an objective function of the form

$$F(Z, \mathcal{D}) = \sum_{(v_i, v_c) \in \mathcal{D}} m((v_i, v_c)) \cdot g_{ic}(Z) \quad (3.1)$$

with respect to the $d \times |V|$ matrix Z . Here, Z is the matrix of vertex embeddings whose columns are the individual

Algorithm 1 Generating Corpus \mathcal{D} on graph $G = (V, E)$

```

1:  $m \leftarrow \mathbf{0}_{|V| \times |V|}$ 
2: for  $n \leftarrow 1$  to  $N$  do
3:    $v_0^n \leftarrow V_0 \sim f_{V_0}$ 
4:    $(v_0^n, v_1^n, \dots, v_{L-1}^n) \leftarrow \text{MARKOV}(v_0, P, L)$ 
5:   for  $j \leftarrow 0$  to  $L - T - 1$  do
6:     for  $r \leftarrow 1$  to  $T$  do
7:        $m(v_j^n, v_{j+r}^n) \leftarrow m(v_j^n, v_{j+r}^n) + 1$ 
8:        $m(v_{j+r}^n, v_j^n) \leftarrow m(v_{j+r}^n, v_j^n) + 1$ 
9:     end for
10:  end for
11: end for
12: return  $\mathcal{D} = (V \times V, m)$ 

```

z_i , d is the dimension of each embedding, and g_{ic} is some $g_{ic} : \mathbb{R}^{d \times |V|} \rightarrow \mathbb{R}$.

In the following, we will use $f = F(\cdot, \mathcal{D} \in \mathbb{D})$ to refer to the *partial application* of an objective function where the corpus has been fixed. This partial application of F yields a function $f : \mathcal{X} \rightarrow \mathbb{R}$ that calculates the objective value for a certain vertex embedding; different corpora now correspond to different f for the same objective function F .

The second stage obtains the vertex embeddings by fixing the corpus \mathcal{D} generated in the first stage from Algorithm 1 in the objective function F to create a partial application $f : \mathcal{X} \rightarrow \mathbb{R}$ which then is optimised with respect to the vertex embeddings by an optimisation procedure A . To keep our discussion as general as possible, we view the optimisation procedure as a deterministic function $A : f \mapsto \mathcal{X}$ which maps an objective function (with the corpus fixed) onto the space of vertex embeddings (the remaining argument of f); stochastic components can be incorporated by fixing the random seed of a pseudo-random number generator.

4 THEORETICAL ANALYSIS

We can now give our main theoretical results. We first prove convergence of the vertex co-occurrences as $N \rightarrow \infty$ and in the joint limit $N, L \rightarrow \infty$ in Section 4.1. We then quantify the rate of these convergences and show how they lead to a heuristic for choosing the hyperparameters N, L in Section 4.2. We finally show in Section 4.3 that, under relatively weak assumptions, convergence of the vertex co-occurrences implies the convergence of the learned vertex representation.

4.1 Convergence of Vertex Co-Occurrences

In this section, we derive the limit of the vertex co-occurrences (the frequencies with which a given vertex pair appears in the corpus) as the number of random walks N and then also their length L become infinite.

Theorem 1 (Occurrence Vertex Co-Occurrences). *Let \mathcal{D} be a corpus generated by Algorithm 1. Retaining the notation from that*

algorithm, the expected relative frequency of occurrence of the pair (v, c) in the corpus is

$$\mathbb{E} \left[\frac{m(v, c)}{|\mathcal{D}|} \right] = \mathbb{E}_{s \sim f_{V_0}} \left[\frac{1}{L-T} \sum_{j=0}^{L-T-1} \frac{1}{2T} \sum_{r=1}^T (P^j)_{sv} (P^r)_{vc} + (P^j)_{sc} (P^r)_{cv} \right].$$

Furthermore,

$$\frac{m(v, c)}{|\mathcal{D}|} - \mathbb{E} \left[\frac{m(v, c)}{|\mathcal{D}|} \right] \rightarrow 0 \text{ almost surely as } N \rightarrow \infty.$$

Proof: A corpus \mathcal{D} generated from N random walks is the union of N independent sub-corpora $\{\mathcal{D}_n\}_{1 \leq n \leq N}$, which are each generated from just a single random walk. Thus the expected value of $\frac{m(v, c)}{|\mathcal{D}|}$ is the expected value of $\frac{m_n(v, c)}{|\mathcal{D}_n|}$ (the relative occurrence frequency in sub-corpus \mathcal{D}_n), averaged over all n . Since the sub-corpora are independent, we have

$$\mathbb{E} \frac{m_n(v, c)}{|\mathcal{D}_n|} = \mathbb{E} \frac{m_m(v, c)}{|\mathcal{D}_m|} \quad \forall n, m \quad (4.1)$$

and hence

$$\mathbb{E} \frac{m(v, c)}{|\mathcal{D}|} = \mathbb{E} \frac{m_n(v, c)}{|\mathcal{D}_n|} \quad \forall n. \quad (4.2)$$

In the remainder of the proof we can therefore focus on a corpus \mathcal{D}_n generated by just a single random walk. We again denote with v_j the j th vertex in this random walk generating \mathcal{D}_n .

The corpus \mathcal{D}_n contains all the pairs (w, c) such that for $0 \leq j \leq L - T - 1$ and $1 \leq r \leq T$ either $(w = v_j, c = v_{j+r})$ or $(w = v_{j+r}, c = v_j)$ (cf. Algorithm 1).

$\mathbb{E} \frac{m_n(w, c)}{|\mathcal{D}_n|}$ is the probability with which we would draw the pair (w, c) from the corpus \mathcal{D}_n . Since \mathcal{D}_n consists of the pairs from just a single (realisation of the) random walk $(v_i)_{0 \leq i \leq L-1}$,

$$\mathbb{P}(\text{draw } (w, c) \text{ from } \mathcal{D}_n) = \mathbb{P}(\text{draw } (w, c) \text{ from } (v_i)_{0 \leq i \leq L-1}). \quad (4.3)$$

The number of pairs (v_i, v_j) in the corpus with $i > j$ is the same as the number of pairs with $i < j$. For a fixed start-vertex s , fixed j , and fixed r , the probability of drawing the vertex pair (w, c) from the random walk therefore is

$$\mathbb{P}(\text{draw } (w, c) \mid s, j, r) = \frac{1}{2} [\mathbb{P}(w = v_j, c = v_{j+r} \mid s, j, r) + \mathbb{P}(w = v_{j+r}, c = v_j \mid s, j, r)]. \quad (4.4)$$

Also,

$$\mathbb{P}(w = v_j, c = v_{j+r} \mid s, j, r) \quad (4.5)$$

$$= \mathbb{P}(w = v_j \mid s = v_0) \cdot \mathbb{P}(c = v_{j+r} \mid w = v_j, s = v_0) \quad (4.6)$$

$$= \mathbb{P}(w = v_j \mid s = v_0) \cdot \mathbb{P}(c = v_{j+r} \mid w = v_j) \quad (4.7)$$

$$= (P^j)_{sw} \cdot (P^r)_{wc} \quad (4.8)$$

where in going from (4.6) to (4.7) we used the Markov Property for the random walk v_i . We similarly obtain

$$\mathbb{P}(w = v_{j+r}, c = v_j \mid s, j, r) = (P^j)_{sc} \cdot (P^r)_{cv}. \quad (4.9)$$

The distributions of j and r are uniform (the number of pairs with some j, r in \mathcal{D}_n is the same for all i, j), so to evaluate the expectation values with respect to j and r we take the arithmetic mean. From Algorithm 1 we see that the distribution of s is f_{V_0} , so the first part of the theorem follows.

For the second part, note that $\frac{m(v, c)}{|\mathcal{D}|}$ is exactly the mean of the independent and identically distribution random variables $\frac{m_n(v, c)}{|\mathcal{D}_n|}$, which take values in $[0, 1]$. The law of large numbers therefore applies to show the almost sure convergence claimed. \square

Corollary 1.1 (Effect of Longer Random Walks). *Using the ergodic theorem (cf. Theorem 1.10.2 in [40]), we can evaluate first the sum over j , and then the expectation value with respect to s , as $L \rightarrow \infty$.*

$$\mathbb{E} \left[\frac{m(v, c)}{|\mathcal{D}|} \right] \xrightarrow{\text{a.s.}} \frac{1}{2T} \sum_{r=1}^T \pi_v(P^r)_{vc} + \pi_c(P^r)_{cv}. \quad (1.1.1)$$

Hence, as $N, L \rightarrow \infty$,

$$\frac{m(v, c)}{|\mathcal{D}|} \xrightarrow{\text{a.s.}} \frac{1}{2T} \sum_{r=1}^T \pi_v(P^r)_{vc} + \pi_c(P^r)_{cv}. \quad (1.1.2)$$

Remark. This almost sure convergence is stronger than the convergence in probability established in [26]. This is due to the fact that they do not let $N \rightarrow \infty$.

Corollary 1.2 (Effect of Starting Distribution). *From the definition of the stationary distribution we see that if $f_{V_0} = \pi$,*

$$\mathbb{E}_{s \sim f_{V_0}} (P^j)_{sv} = \pi_v. \quad (1.2.3)$$

In that case, (1.1.1) holds as an equality for arbitrary $L > T$, and (1.1.2) holds for arbitrary fixed L as $N \rightarrow \infty$.

The results from this section, together with Theorem 7, guarantee that if the DeepWalk algorithm performs enough and long enough random walks, the learned vertex representations converge in the sense of Theorem 7.

4.2 Convergence Rates

Let us first investigate the convergence rate as $N \rightarrow \infty$ for arbitrary L .

Theorem 2 (Convergence Rate as $N \rightarrow \infty$). *The convergence of Theorem 1 satisfies the concentration bound*

$$\mathbb{P} \left(\left| \frac{m(v, c)}{|\mathcal{D}|} - \mathbb{E} \left[\frac{m(v, c)}{|\mathcal{D}|} \right] \right| > \varepsilon \right) < 2 \exp(-2N\varepsilon^2) \quad (2.1)$$

for $\varepsilon > 0$.

Proof: Let again $m_n(v, c)$ be the multiplicity function of the corpus \mathcal{D}_n generated from the n th random walk. Each

random walk has the same length, so $N \cdot |\mathcal{D}_n| = |\mathcal{D}| \forall n < N$. Since $0 \leq m_n(v, c) \leq |\mathcal{D}_n|$, Hoeffding's inequality gives

$$\mathbb{P} \left(\left| \sum_{n=1}^N m_n(v, c) - \mathbb{E} \left[\sum_{n=1}^N m_n(v, c) \right] \right| > \varepsilon \right) < 2 \exp \left(-\frac{2\varepsilon^2}{N|\mathcal{D}_n|^2} \right). \quad (4.10)$$

Dividing by $|\mathcal{D}| = N|\mathcal{D}_n|$ and using $\sum_{n=1}^N m_n(v, c) = m(v, c)$ then yields

$$\mathbb{P} \left(\left| \frac{m(v, c)}{|\mathcal{D}|} - \mathbb{E} \left[\frac{m(v, c)}{|\mathcal{D}|} \right] \right| > \frac{\varepsilon}{N|\mathcal{D}_n|} \right) < 2 \exp \left(-\frac{2\varepsilon^2}{N|\mathcal{D}_n|^2} \right). \quad (4.11)$$

Re-scaling $\varepsilon \rightarrow N|\mathcal{D}_n|\varepsilon$ then yields the theorem. \square

To derive a concentration bound for the joint limit, we first need to prove two lemmas that quantify the convergence rate of $\mathbb{E} \left[\frac{m(v, c)}{|\mathcal{D}|} \right]$ as $L \rightarrow \infty$ for undirected and directed graphs. For notational simplicity, define

$$\omega(v, c) := \frac{1}{2T} \sum_{r=1}^T \pi_v(P^r)_{vc} + \pi_c(P^r)_{cv}$$

as the value that $\mathbb{E} \left[\frac{m(v, c)}{|\mathcal{D}|} \right]$ converges to (Theorem 1).

Lemma 3 (Convergence Rate for Undirected Graphs). *Given an undirected graph, let $d_i = \sum_j w_{ij}$ be the out-degrees and $D = \text{diag}(d_i)$. Let $\{\lambda_k\}$ be the ordered eigenvalues of the symmetric normalised graph Laplacian*

$$L^{\text{norm}} = I - D^{1/2} P D^{-1/2}.$$

Then we have the convergence estimate

$$\begin{aligned} & \left| \mathbb{E} \left[\frac{m(v, c)}{|\mathcal{D}|} \right] - \omega(v, c) \right| \\ & \leq \frac{1}{L-T} \frac{1 - \phi_*^{L-T}}{1 - \phi_*} \mathbb{E}_{s \sim f_{V_0}} \left[d_s^{-\frac{1}{2}} \right] \\ & \quad \cdot \frac{1}{2T} \sum_{r=1}^T \sqrt{d_v(P^r)_{vc}} + \sqrt{d_c(P^r)_{cv}} = O \left(\frac{1}{L} \right) \end{aligned}$$

where $\phi_* \equiv \mu_* = \sup\{|1 - \lambda_2|, |1 - \lambda_{|V|}|\}$ for non-bipartite graphs and $\phi_* \equiv \nu_* = \sup\{|1 - \lambda_2|, |1 - \lambda_{|V|-1}|\}$ for bipartite graphs and even $L - T$.

Proof: For a general undirected graph, we get by substituting for $\omega(v, c)$ and from Theorem 1 for $\mathbb{E} \left[\frac{m(v, c)}{|\mathcal{D}|} \right]$

$$\begin{aligned} & \left| \mathbb{E} \left[\frac{m(v, c)}{|\mathcal{D}|} \right] - \omega(v, c) \right| \\ & = \frac{1}{2T} \mathbb{E}_{s \sim f_{V_0}} \mathbb{E}_j \sum_{r=1}^T [(P^j)_{sv} - \pi_v](P^r)_{vc} \\ & \quad + [(P^j)_{sc} - \pi_c](P^r)_{cv} \quad (4.12) \end{aligned}$$

$$\begin{aligned} & \leq \frac{1}{2T} \mathbb{E}_{s \sim f_{V_0}} \sum_{r=1}^T \left[|\mathbb{E}_j (P^j)_{sv} - \pi_v| (P^r)_{vc} \right. \\ & \quad \left. + |\mathbb{E}_j (P^j)_{sc} - \pi_c| (P^r)_{cv} \right] \quad (4.13) \end{aligned}$$

by Jensen's inequality. For non-bipartite graphs, we again apply Jensen's inequality and use Theorem 5.1 in [41] to obtain

$$\begin{aligned} \text{LHS} & \leq \frac{1}{2T} \mathbb{E}_{s \sim f_{V_0}} \mathbb{E}_j \sum_{r=1}^T \sqrt{\frac{d_v}{d_s}} \mu_*^j (P^r)_{vc} \\ & \quad + \sqrt{\frac{d_c}{d_s}} \mu_*^j (P^r)_{cv} \quad (4.14) \end{aligned}$$

$$\begin{aligned} & \leq \frac{1}{L-T} \frac{1 - \mu_*^{L-T}}{1 - \mu_*} \mathbb{E}_{s \sim f_{V_0}} \left[d_s^{-\frac{1}{2}} \right] \\ & \quad \cdot \frac{1}{2T} \sum_{r=1}^T \sqrt{d_v(P^r)_{vc}} + \sqrt{d_c(P^r)_{cv}} \quad (4.15) \end{aligned}$$

For $L - T$ even (so that \mathbb{E}_j contains an even number of terms) for bipartite graphs, we replace μ_* with ν_* , where the difference between the two stems from the fact that $1 - \lambda_{|V|} = -1$, so in the derivation of Theorem 5.1 in [41] successive terms involving $(1 - \lambda_{|V|})^t$ cancel. The Lemma then follows. \square

Lemma 4 (Convergence Rate for Directed Graphs). *For a directed graph G with period $\Theta \geq 1$,*

$$\left| \mathbb{E} \left[\frac{m(v, c)}{|\mathcal{D}|} \right] - \omega(v, c) \right| = O \left(\frac{1}{L} \right). \quad (4.1)$$

Furthermore, if G is aperiodic ($\Theta = 1$), with constants C and $\alpha \in (0, 1)$ as in Theorem 4.9 in [42],

$$\begin{aligned} & \left| \mathbb{E} \left[\frac{m(v, c)}{|\mathcal{D}|} \right] - \omega(v, c) \right| \\ & \leq \frac{C}{L-T} \frac{1 - \alpha^{L-T}}{1 - \alpha} \frac{1}{2T} \sum_{r=1}^T (P^r)_{vc} + (P^r)_{cv}. \quad (4.2) \end{aligned}$$

To prove this for both aperiodic and periodic graphs, we need the following intermediate result.

Proposition 5. *Let P be the transition matrix of a Markov chain $\{X_t\}_{t \geq 0}$ with period Θ and stationary distribution π . Let $\{\mathcal{S}_i\}_{1 \leq i \leq \Theta}$ be the disjoint subsets of the state space S such that $\mathbb{P}(X_{t+1} \in \mathcal{S}_{(i+1) \bmod \Theta} | X_t \in \mathcal{S}_i) > 0$. If vertices $s, v \in \mathcal{S}_i$ for some i , then there exist (explicitable) constants $\alpha \in (0, 1)$, $C > 0$ such that*

$$\left| (P^{\Theta \cdot t})_{sv} - \Theta \cdot \pi_v \right| \leq C \alpha^t.$$

Proof: A similar observation is well-known for Markov chains in general (Theorem 4.9, [42]); in the present context, it suffices to observe that $P^\Theta|_{\mathcal{S}_i}$ is an aperiodic stochastic matrix on the block \mathcal{S}_i , for which $\Theta \pi|_{\mathcal{S}_i}$ is an invariant probability measure, and the lemma follows from the cited result. \square

With this proposition in hand, we are now in a position to prove the Lemma.

Proof of Lemma 4: Substituting for $\omega(v, c)$ and from Theorem 1 for $\mathbb{E} \left[\frac{m(v, c)}{|\mathcal{D}|} \right]$ and using Jensen's inequality as in the proof for Lemma 3 again leads to (4.13).

We prove the first part of the theorem by showing that for a general periodic graph $|\mathbb{E}_j (P^j)_{sv} - \pi_v| = O \left(\frac{1}{L} \right)$.

Let again $\{\mathcal{S}_i\}_{1 \leq i \leq \Theta}$ be disjoint subsets of the state space S as in Proposition 5. Without loss of generality, take $s \in \mathcal{S}_i$ and $v \in \mathcal{S}_k$ with $k \geq i$, and let $a = k - i$ so that $\sum_{s' \in \mathcal{S}_k} (P^a)_{ss'} = 1$. We decompose the path from s to v in j steps, $(P^j)_{sv}$, into a part from $s \in \mathcal{S}_i$ to an $s' \in \mathcal{S}_k$ and a part from $s' \in \mathcal{S}_k$ to $v \in \mathcal{S}_k$:

$$\mathbb{E}_j(P^j)_{sv} = \frac{1}{L-T} \sum_{j=0}^{L-T-1} (P^j)_{sv} \quad (4.16)$$

$$= \frac{1}{L-T} \sum_{j=a}^{L-T-1} \sum_{s' \in \mathcal{S}_k} (P^a)_{ss'} (P^{j-a})_{s'v} \quad (4.17)$$

$$= \frac{1}{L-T} \sum_{n=0}^{\lfloor \frac{L-T-a-1}{\Theta} \rfloor} (P^{\Theta \cdot n})_{s'v} \quad (4.18)$$

since if $j < a$ then $(P^j)_{sv} = 0$, and if $(j-a) \bmod \Theta \neq 0$ then $(P^{j-a})_{s'v} = 0$, and since $\sum_{s' \in \mathcal{S}_k} (P^a)_{ss'} = 1$. We can use Proposition 5 to obtain upper and lower bounds for $(P^{\Theta \cdot n})_{s'v}$

$$\Theta \cdot \pi_v - C\alpha^n \leq (P^{\Theta \cdot n})_{s'v} \leq \Theta \cdot \pi_v + C\alpha^n \quad (4.19)$$

for some constants C and $\alpha \in (0, 1)$. If $\mathbb{E}_j(P^j)_{sv} - \pi_v \geq 0$, we use the upper bound on $(P^{\Theta \cdot n})_{s'v}$, and if $\mathbb{E}_j(P^j)_{sv} - \pi_v \leq 0$ we use the lower bound to obtain:

$$\begin{aligned} \text{LHS} &\leq \left| \frac{1}{L-T} \sum_{n=0}^{\lfloor \frac{L-T-a-1}{\Theta} \rfloor} (\Theta \cdot \pi_v) - \pi_v \right| \\ &\quad + \frac{1}{L-T} \sum_{n=0}^{\lfloor \frac{L-T-a-1}{\Theta} \rfloor} (C\alpha^n). \end{aligned} \quad (4.20)$$

Using the fact that $\lfloor \frac{x}{y} \rfloor = \frac{x}{y} - \frac{x \bmod y}{y}$ to evaluate the prefactor of π_v , we continue

$$\begin{aligned} \text{LHS} &\leq \left| \frac{\Theta}{L-T} \left(\left\lfloor \frac{L-T-a-1}{\Theta} \right\rfloor + 1 \right) - 1 \right| \pi_v \\ &\quad + \frac{1}{L-T} \sum_{n=0}^{\lfloor \frac{L-T-a-1}{\Theta} \rfloor} C\alpha^n \end{aligned} \quad (4.21)$$

$$\leq \left| \frac{\Theta - a - 1 - (L-T-a-1) \bmod \Theta}{L-T} \right| \cdot \pi_v$$

$$+ \frac{1}{L-T} \sum_{n=0}^{\lfloor \frac{L-T-a-1}{\Theta} \rfloor} C\alpha^n \quad (4.22)$$

$$= O\left(\frac{1}{L}\right) \quad (4.23)$$

since the remaining geometric sum is bounded by a constant.

To prove the second part of the theorem, note that for aperiodic graphs, $\Theta = 1$ and $i = k \Rightarrow a = 0$. For aperiodic graphs, C and α still have the same meaning as before.

Therefore, in (4.22) the prefactor of π_v vanishes for aperiodic graphs. The upper bound on $|\mathbb{E}_j(P^j)_{sv} - \pi_v|$ therefore evaluates to

$$|\mathbb{E}_j(P^j)_{sv} - \pi_v| \leq \frac{C}{L-T} \sum_{n=0}^{L-T-1} \alpha^n \quad (4.24)$$

$$\leq \frac{C}{L-T} \frac{1 - \alpha^{L-T}}{1 - \alpha}. \quad (4.25)$$

Substituting into (4.13), the theorem follows. \square

Finally we state and prove our theorem for the convergence rate in the joint limit $N, L \rightarrow \infty$:

Theorem 6 (Convergence Rate as $N, L \rightarrow \infty$). *The convergence of Corollary 1.1 satisfies the concentration bound*

$$\mathbb{P} \left(\left| \frac{m(v, c)}{|\mathcal{D}|} - \omega(v, c) \right| > \varepsilon \right) < 2 \exp(-2N(\varepsilon - U)^2) \quad (6.1)$$

for $\varepsilon > U > 0$, and where U is the appropriate upper bound for $\left| \mathbb{E} \left[\frac{m(v, c)}{|\mathcal{D}|} \right] - \omega(v, c) \right|$ as derived in lemmas 3 and 4 and depends only on L, T , and graph-specific constants.

Proof: We have by the triangle inequality

$$\begin{aligned} &\left| \frac{m(v, c)}{|\mathcal{D}|} - \mathbb{E} \left[\frac{m(v, c)}{|\mathcal{D}|} \right] \right| \\ &\geq \left| \frac{m(v, c)}{|\mathcal{D}|} - \omega(v, c) \right| - \left| \mathbb{E} \left[\frac{m(v, c)}{|\mathcal{D}|} \right] - \omega(v, c) \right|. \end{aligned}$$

Therefore,

$$\begin{aligned} &\mathbb{P} \left(\left| \frac{m(v, c)}{|\mathcal{D}|} - \mathbb{E} \left[\frac{m(v, c)}{|\mathcal{D}|} \right] \right| > \varepsilon \right) \\ &\geq \mathbb{P} \left(\left| \frac{m(v, c)}{|\mathcal{D}|} - \omega(v, c) \right| - \left| \mathbb{E} \left[\frac{m(v, c)}{|\mathcal{D}|} \right] - \omega(v, c) \right| > \varepsilon \right) \\ &= \mathbb{P} \left(\left| \frac{m(v, c)}{|\mathcal{D}|} - \omega(v, c) \right| > \varepsilon + \left| \mathbb{E} \left[\frac{m(v, c)}{|\mathcal{D}|} \right] - \omega(v, c) \right| \right) \\ &\geq \mathbb{P} \left(\left| \frac{m(v, c)}{|\mathcal{D}|} - \omega(v, c) \right| > \varepsilon + U \right). \end{aligned}$$

Substituting $\varepsilon \rightarrow \varepsilon - U$ and using Theorem 2 then yields the result. \square

We can use Theorem 6 to obtain a rough heuristic for an "optimal" choice of the parameters N and L given a fixed computational cost; the window size T determines the size of the local structures which the vertex embeddings encode and is therefore usually fixed for a given application. The total computational cost of computing the co-occurrence matrix scales linearly with the number of random-walk steps $N \cdot L$, but since $L \geq T$ we define the *excess computational cost* $K = N \cdot L - N \cdot T = N \cdot (L - T)$.

For notational simplicity we also denote the upper bound for the failure probability in Theorem 6 as $\delta = 2 \exp(-2N(\varepsilon - U)^2)$ and absorb all terms in U that are independent of N, L into a constant g for notational simplicity. Then we have:

Corollary 6.1. *Assuming that $K \gg N$ so that $\phi^{K/N} \approx 0$ and $\alpha^{K/N} \approx 0$ in U , the value for N that minimises the error $\varepsilon = U + \sqrt{\frac{1}{2N} \log \frac{2}{\delta}}$ given K, δ fixed is*

$$N \approx \frac{1}{2} \sqrt[3]{\frac{K^2}{g^2} \log \frac{2}{\delta}}. \quad (6.1.2)$$

This heuristic minimises the error ε given an excess computational cost K and an upper bound for the failure probability δ , which both can be freely chosen. We see that under the heuristic, $N = O\left(K^{\frac{2}{3}}\right)$, which justifies the assumption $K \gg N$. Even without this assumption,

we still have an upper bound, as the neglected terms $-\phi^{K/N}$, $-\alpha^{K/N}$ can only make ε smaller. In practice, one could use $g \approx 1$ as a first approximation to the remaining constant g , which depends on the graph structure through the spectrum or through α . If a more precise estimate for g is needed, let us remark that this proof (and the results recalled in the Appendix) produce a constructable g , in that one can follow the steps of the proofs to explicitly find g for a given example. In the case where we use Lemma 3, let us remark that many techniques are known ([43], and references therein) for bounding spectral gaps, which would lead to a constructable upper bound for g .

4.3 Convergence of Vertex Representations

We finally show that convergence of the corpora, suitably normalised, is indeed sufficient to show convergence of the vertex representations.

We first remark that, if two corpora $\mathcal{D}^i = (D^{(i)}, m^{(i)})$, $i = 1, 2$ differ only by a constant multiple so that

$$D^{(1)} = D^{(2)}; \quad m^{(1)} = c \cdot m^{(2)}$$

for some $c > 0$, then the respective partial applications of the objective function also differ only up to a constant multiple: $f^{(1)} = c f^{(2)}$, and it follows that the derived vertex embeddings

$$\{z_i^{(1)}\}_{i: v_i \in V} = \{z_i^{(2)}\}_{i: v_i \in V} \quad (4.26)$$

also agree. To remove this redundancy, we define normalised objective functions

$$f^{\text{norm}} = \frac{1}{|\mathcal{D}|} \cdot f. \quad (4.27)$$

By (4.26) this normalisation does not change the learned representations.

We next assume the following:

Assumption 1. The objective function is of the same form as in (3.1), with the g_{ic} bounded.

Assumption 2. The map

$$f \mapsto f(A(f)) \text{ is continuous} \quad (4.28)$$

on some space of partial applications of objective functions, with respect to the essential supremum distance $\|\cdot\|_\infty$, where A is the optimisation procedure $A: f \mapsto \mathcal{X}$ as before.

Informally, Assumption 2 means that if the partial application of the objective function, f , changes slightly, the optimisation procedure will find vertex embeddings that have a similar objective value.

Theorem 7 (Convergence of Vertex Representations). *Let (\mathcal{D}_n) be a sequence of corpora generated by Algorithm 1 from the same graph such that the respective vertex pair occurrence frequencies $\left(\frac{m_n}{|\mathcal{D}_n|}\right)$ converge to the vertex pair occurrence frequencies $\frac{m}{|\mathcal{D}|}$ in some corpus \mathcal{D} generated from that graph. Let f_n be the partial application of an objective function given corpus \mathcal{D}_n , and f_n^{norm} the normalised objective function, and similarly define f and f^{norm} . Then*

$$f_n^{\text{norm}}(A(f_n)) \longrightarrow f^{\text{norm}}(A(f)).$$

Table 1: Characteristics of the datasets used in our experiments.

DATASET	# OF Classes	# OF NODES V	# OF EDGES E
CORA	7	2,708	5,429
BLOGCATALOG	39	10,312	333,983
WIKI	19	2,405	17,981
FACEBOOK LARGE	4	22,470	171,002
YOUTUBE	47	1.1M	4.9M




Fig. 2. Graph visual representation of some of the datasets used in our experimental results.

Proof: The map $\frac{m(\cdot)}{\mathcal{D}} \rightarrow f^{\text{norm}}$ is continuous with respect to the uniform distance on functions, as the g_{ic} are bounded.

By continuity of f_n^{norm} with respect to $\left(\frac{m_n}{|\mathcal{D}_n|}\right)$, $\left(\frac{m_n}{|\mathcal{D}_n|}\right) \rightarrow \frac{m}{|\mathcal{D}|}$ implies $f_n^{\text{norm}} \rightarrow f^{\text{norm}}$.

By continuity of $f_n^{\text{norm}}(A(f_n^{\text{norm}}))$ with respect to f_n^{norm} , $f_n^{\text{norm}} \rightarrow f^{\text{norm}}$ implies $f_n^{\text{norm}}(A(f_n^{\text{norm}})) \rightarrow f^{\text{norm}}(A(f^{\text{norm}}))$.

From (4.26) we know that $A(f_n^{\text{norm}}) = A(f_n)$ and $A(f^{\text{norm}}) = A(f)$ and the theorem follows. \square

Theorem 7 means that if the vertex pair occurrence frequencies of (vertex, context) pairs converge to those of some corpus \mathcal{D} , then the vertex representations converge to values that have the same (normalised) objective function value as the representations derived from corpus \mathcal{D} and should be similarly useful for downstream tasks.

5 EXPERIMENTAL RESULTS

In this section, we detail the set of experiments that we conducted to illustrate our theoretical findings.

5.1 Datasets Description

We used five real-world datasets for our experiments. Each is a graph whose vertices are grouped in a number of classes. They vary in their number of nodes and edges (up to the order of millions) and in their number of classes. The dataset statistics are given in Table 1 along with a graphical visualisation (see Fig. 2).

- The BLOGCATALOG [44] dataset is derived from a blog-sharing website. The edges reflect the bloggers (vertices) following each other. The classes in the dataset are the bloggers' interests.
- The CORA [45] dataset is a text classification dataset for Machine Learning papers. Each vertex corresponds to a paper with the edges representing citation links. The classes are different areas in Machine Learning. The dataset also supplies bag-of-word

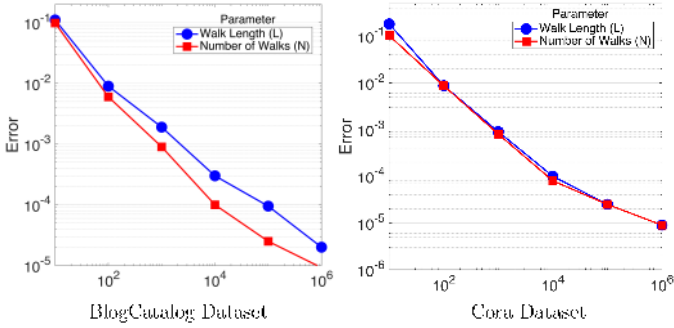


Fig. 3. Convergence of co-occurrence matrices as the trajectory length L and the number of walks N approach infinity. The error is the computed as $\|m - \omega\|_2$, and we use $N = 80$ when varying L and $L = 40$ when varying N .

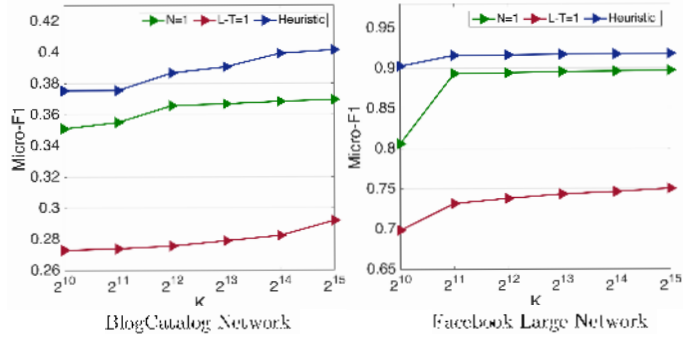


Fig. 4. Computational cost in the x-axis expressed as $K = N \cdot (L - T)$ against the task performance, in terms of Micro-F1, on the y-axis. The plots compare our derived heuristic with two other cases where the parameters are not optimal.

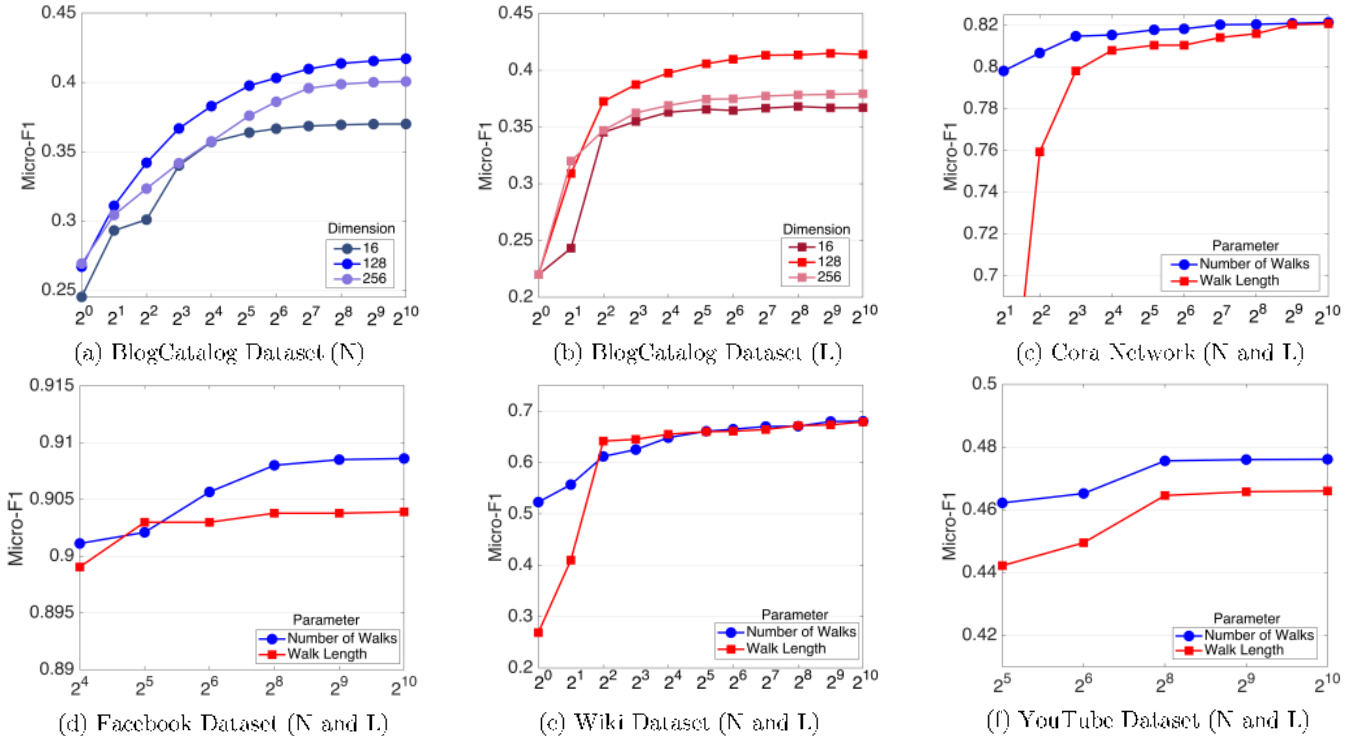


Fig. 5. Performance on a classification task for different using embeddings generated with different hyperparameter choices for the number of random walks N , the length of each random walk L , and the embedding dimension d . We use $N = 80$ when varying L and $L = 40$ when varying N , and fix $d = 128$ for plots (c)-(f).

representations for each vertex, though we did not use these in our experiments.

- The WIKI [17] dataset, similarly to Cora, is a collection of text documents (vertices) with edges representing hyperlinks between them.
- The FACEBOOK LARGE [46] dataset is a page-page graph of Facebook sites. The vertices refer to Facebook pages while the edges denote mutual likes between sites.
- The YOUTUBE [47] dataset is based on the online video-sharing platform, with edges reflecting subscription connections between users (vertices).

5.2 Results & Discussion

We support our theoretical analysis following a three-part evaluation scheme.

(E1) Convergence of co-occurrence matrices. We begin by directly validating our results about the speed of convergence for the corpus generated by the first stage of DeepWalk (Theorems 2 and 6).

To do this, we calculate the approximation error of the co-occurrence matrix m as the Frobenius norm of the difference between m and its asymptotic value ω , i.e. $\|m - \omega\|_2$, for different values of the random walk length L and the number of random walks N .

More precisely, for all experiments run in this part we

set the window size $T = 10$ and the embedding dimension $d = 128$. To observe the effect of increasing the number of walks, we set $N = \{10^1, 10^2, 10^3, 10^4, 10^5, 10^6\}$ and fix $L = 40$ (see blue lines in Fig. 3). The effect of increasing the length of the walk, we set $L = \{10^1, 10^2, 10^3, 10^4, 10^5, 10^6\}$ and fix $N = 80$ (see red lines in Fig. 3). We ran these experiments using the BLOGCATALOG and CORA datasets. The results are displayed in Fig. 3, where one can observe a convergence rate consistent with our results for the limiting cases when $N \rightarrow \infty$ and $L \rightarrow \infty$.

(E2) Heuristic for optimal values for N and L . We show the practical application of our heuristic for optimal choices of the hyperparameters N and L for a fixed computational complexity $K = N \cdot (L - T)$ and fixed failure probability δ (Corollary 6.1). For our experiments, we plot the task performance against the computational cost K when choosing N and L using three different strategies. In the first strategy, we set $L = T + 1$ and $N = K$ to simulate the extreme of the minimum random walk length (red line in Fig. 4). In the second strategy, we set $N = 1$ and $L = T + K$, the case of performing a single very long walk, which corresponds to the hyperparameter choice suggested in [37] (green line in Fig. 4). Finally, we select N and L using our heuristic, approximating the constant $g \approx 1$ and fixing $\delta = 0.01$ (blue line in Fig. 4).

The task performance is the Micro-F1 score achieved by the one-vs-rest logistic regression suggested by [18] and also used in [26], [37]. For all strategies we again set $T = 10$ and $d = 128$. We ran these experiments using the BLOGCATALOG and the FACEBOOK LARGE datasets for the values $K = \{2^{10}, 2^{11}, 2^{12}, 2^{13}, 2^{14}, 2^{15}\}$, so after necessary rounding the values used to test the heuristic were $N = \{88, 140, 224, 355, 563, 892\}$ and $L - T = \{12, 15, 18, 23, 29, 37\}$. The results are plotted in Fig. 4.

Choosing N and L using the heuristic consistently outperforms the other two cases, illustrating the impact that a better choice of the hyperparameters N and L can have. The maximally-parallelisable choice of $N = K$ and $L = T + 1$ performs a lot worse than the other two strategies, which is likely due to the uniform starting distribution f_{V_0} being too different to the stationary distribution on these graphs, resulting in the vertex contexts generated from the short walks not being very informative.

(E3) Convergence of task performance. In practice, vertex embeddings are seldom the final goal of a machine learning pipeline, and our notion of convergence in Theorem 7 takes even two different sets of vertex embeddings to be equivalent if they result in the same (normalised) objective function value. We therefore use convergence of the task performance of a down-stream algorithm as a proxy for the convergence of the vertex embeddings; if the downstream task performance has converged then the vertex embeddings have very likely converged as well (in the sense of Theorem 7).

We again use the standard protocol for network embeddings evaluation ([18], [26], [37]) and compute the Micro-F1 score of a one-vs-rest logistic regression. We again set the window size $T = 10$. The results are plotted in Fig. 5. For the experiments in Fig. 5 (a) and (b) we use varying embedding dimensions of $d = \{16, 128, 256\}$, while we keep the embedding dimension constant at $d = 128$ for the other

experiments in plots (c)-(f). When varying L , we set $N = 80$, and when varying N , we set $L = 40$.

One observes that as the number of walks and the length of each walk are tending to large values the performance converges. This observation holds across all graph sizes and dimensions, for both medium-sized graphs like the CORA dataset and large graphs like the BLOGCATALOG dataset and embedding dimensions ranging from $d = 16$ to $d = 256$.

6 CONCLUSION

We presented a convergence analysis for random-walk-based vertex embeddings. Unlike existing theoretical works, we investigated the convergence of vertex co-occurrences in the generated corpus as the number of random walks $N \rightarrow \infty$ for arbitrary random walk lengths L as well as the simultaneous limit $N, L \rightarrow \infty$, proving almost sure convergence and quantifying the convergence rates by deriving concentration bounds for the two limits. We also proved that convergence of vertex co-occurrences does indeed imply convergence of the vertex embeddings themselves, a fact that to our knowledge had not previously received formal theoretical treatment. Moreover, we provided a heuristic for choosing optimal values for the hyperparameters N and L . We supported our theory with a set of experiments using several real-world graphs, and demonstrate the practical implication of our findings.

APPENDIX

EXTENDED MATHEMATICAL RESULTS

This appendix explicitly gives the mathematical results and some proofs cited, following the notation used in the main paper.

Theorem 8 (Ergodic Theorem; Theorem 1.10.2 in [40]). *Let P be an irreducible transition matrix and let λ be any distribution on the state space V . If $(X_t)_{t \geq 0}$ is a Markov chain with transition matrix P , starting distribution λ , and invariant distribution π , then*

$$\frac{N_i(n)}{n} \rightarrow \frac{1}{m_i} \text{ almost surely as } n \rightarrow \infty \quad (8.1)$$

where $m_i = \mathbb{E}(T_i) = \frac{1}{\pi_i}$ is the expected return time to state $i \in V$ and N_i is the number of visits to state i before time n . Moreover, if P is finite then for any bounded function $f : V \rightarrow \mathbb{R}$ we have

$$\frac{1}{n} \sum_{t=0}^{n-1} f(X_t) \rightarrow \bar{f} = \sum_{i \in S} \pi_i f_i \text{ almost surely as } n \rightarrow \infty.$$

Proof: If P is transient, then, with probability 1, the total number of visits N_i to state i is finite, so

$$\frac{N_i(n)}{n} \leq \frac{N_i}{n} \rightarrow 0 = \frac{1}{m_i}. \quad (A.1)$$

Suppose then that P is recurrent and fix a state i . For $T = T_i$ the return time to state i we then have $\mathbb{P}(T < \infty) = 1$ and $(X_{T+n})_{n \geq 0}$ is a Markov chain with starting distribution δ_i and independent of X_0, X_1, \dots, X_T by the strong Markov property. The long-run proportion of time spent in i is the same for $(X_{T+n})_{n \geq 0}$ and for $(X_n)_{n \geq 0}$, so it suffices to consider the case $\lambda = \delta_i$.

Write $S_i^{(r)}$ as the length of the r th excursion to i , i.e. $S_i^{(r)} = T_i^{(r)} - T_i^{(r-1)}$ where $T_i^{(r)}$ is the r th return time to i . The non-negative random variables $S_i^{(1)}, S_i^{(2)}, \dots$ are independent and identically distributed with $\mathbb{E}(S_i^{(r)}) = m_i$ by the strong Markov property. Now

$$S_i^{(1)} + \dots + S_i^{(N_i(n)-1)} \leq n-1, \quad (\text{A.2})$$

the left-hand-side being the time of the last visit to i before time n . Also,

$$S_i^{(1)} + \dots + S_i^{(N_i(n))} \geq n, \quad (\text{A.3})$$

the left-hand side being the time of the first visit to i after time $n-1$. Hence,

$$\frac{S_i^{(1)} + \dots + S_i^{(N_i(n)-1)}}{N_i(n)} \leq \frac{n}{N_i(n)} \quad (\text{A.4})$$

and

$$\frac{n}{N_i(n)} \leq \frac{S_i^{(1)} + \dots + S_i^{(N_i(n))}}{N_i(n)}. \quad (\text{A.5})$$

By the strong law of large numbers

$$\frac{S_i^{(1)} + \dots + S_i^{(n)}}{n} \rightarrow m_i \text{ almost surely as } n \rightarrow \infty \quad (\text{A.6})$$

and, since P is recurrent,

$$N_i(n) \rightarrow \infty \text{ almost surely as } n \rightarrow \infty. \quad (\text{A.7})$$

So, letting $n \rightarrow \infty$ in (A.4) and (A.5), we get

$$\frac{n}{N_i(n)} \rightarrow m_i \text{ almost surely as } n \rightarrow \infty \quad (\text{A.8})$$

which implies

$$\frac{N_i(n)}{n} \rightarrow \frac{1}{m_i} \text{ almost surely as } n \rightarrow \infty. \quad (\text{A.9})$$

Let $f : V \rightarrow \mathbb{R}$ now be a bounded function and assume without loss of generality that $|f| \leq 1$. For any $W \subseteq V$ we have

$$\begin{aligned} \left| \frac{1}{n} \sum_{k=0}^{n-1} f(X_k) - \bar{f} \right| &= \left| \sum_{i \in V} \left(\frac{N_i(n)}{n} - \pi_i \right) f_i \right| \\ &\leq \sum_{i \in W} \left| \frac{N_i(n)}{n} - \pi_i \right| + \sum_{i \notin W} \left| \frac{N_i(n)}{n} - \pi_i \right| \\ &\leq \sum_{i \in W} \left| \frac{N_i(n)}{n} - \pi_i \right| + \sum_{i \notin W} \left(\frac{N_i(n)}{n} + \pi_i \right) \\ &\leq 2 \sum_{i \in W} \left| \frac{N_i(n)}{n} - \pi_i \right| + 2 \sum_{i \notin W} \pi_i. \end{aligned} \quad (\text{A.10})$$

Given $\varepsilon > 0$, choose W finite such that

$$\sum_{i \notin W} \pi_i < \varepsilon/4. \quad (\text{A.11})$$

We proved above that $\frac{N_i(n)}{n} \rightarrow \pi_i$ almost surely, so choose an L such that for $n \geq L$

$$\sum_{i \in W} \left| \frac{N_i(n)}{n} - \pi_i \right| < \varepsilon/4. \quad (\text{A.12})$$

Then, for $n \geq L$ we have

$$\left| \frac{1}{n} \sum_{k=0}^{n-1} f(X_k) - \bar{f} \right| < \varepsilon, \quad (\text{A.13})$$

which establishes the desired convergence. \square

Proposition 9 (Equation 3.1 in [41]). *Consider an undirected graph with symmetric normalised Laplacian $L^{\text{norm}} = I - D^{1/2} P D^{-1/2}$, transition matrix P , and stationary distribution π . Let the ordered eigenvalues $\{\lambda_k\}$ of L^{norm} have normalised eigenvectors $\{v^{(k)}\}$. Then*

$$(P^t)_{ij} = \pi_j + \sum_{k=2}^{|V|} (1 - \lambda_k)^t v_i^{(k)} v_j^{(k)} \sqrt{\frac{d_j}{d_i}}.$$

Proof: The eigenvectors of a symmetric matrix are mutually orthogonal. We can thus write in spectral form

$$I - L^{\text{norm}} = \sum_{k=1}^{|V|} (1 - \lambda_k) v^{(k)} v^{(k)T}$$

From the definition of L^{norm} , and using the fact that the eigenvector $v^{(1)}$ with eigenvalue $\lambda_1 = 0$ is $v_i^{(1)} = \sqrt{\pi_i}$, we get

$$\begin{aligned} P^t &= D^{-\frac{1}{2}} (I - L^{\text{norm}})^t D^{\frac{1}{2}} \\ &= \sum_{k=1}^{|V|} (1 - \lambda_k)^t D^{-\frac{1}{2}} v^{(k)} v^{(k)T} D^{\frac{1}{2}} \\ &= Q + \sum_{k=2}^{|V|} (1 - \lambda_k)^t D^{-\frac{1}{2}} v^{(k)} v^{(k)T} D^{\frac{1}{2}}. \end{aligned}$$

Using $\pi_i = d_i / \sum_{k \in V} d_k$ we have

$$\begin{aligned} Q_{ij} &= (1 - \lambda_1) (D^{-\frac{1}{2}} v^{(1)} v^{(1)T} D^{\frac{1}{2}})_{ij} \\ &= \frac{1}{\sqrt{d_i}} \sqrt{\pi_i} \sqrt{\pi_j} \sqrt{d_j} \\ &= \pi_j. \end{aligned} \quad (\text{A.14})$$

The proposition follows. \square

Theorem 10 (Mixing Rate of Random Walk on Undirected Graph; Theorem 5.1 in [41]). *For a random walk on an undirected graph, with $\mu_\star = \sup\{|1 - \lambda_2|, |1 - \lambda_{|V|}|\}$ (again, $\{\lambda_k\}$ are the ordered eigenvalues of L^{norm}),*

$$|(P^t)_{ij} - \pi_j| \leq \sqrt{\frac{d_j}{d_i}} \mu_\star^t.$$

Proof: Starting from Proposition 9, we have

$$\begin{aligned} |(P^t)_{ij} - \pi_j| &= \left| \sum_{k=2}^{|V|} (1 - \lambda_k)^t v_i^{(k)} v_j^{(k)} \sqrt{\frac{d_j}{d_i}} \right| \\ &\leq \sqrt{\frac{d_j}{d_i}} \mu_\star^t \cdot \left| \sum_{k=2}^{|V|} v_i^{(k)} v_j^{(k)} \right| \\ &\leq \sqrt{\frac{d_j}{d_i}} \mu_\star^t \end{aligned}$$

by normalisation of the $v^{(k)}$. \square

Corollary 10.1 (Mixing Rate of Random Walk on Bipartite Graph). *For a random walk on a bipartite undirected graph, with $\nu_\star = \sup\{|1 - \lambda_2|, |1 - \lambda_{|V|-1}|\}$, denoting with $\langle \cdot \rangle_t$ an average over an even number of timesteps ($\langle f(t) \rangle_t = \frac{1}{n+1} \sum_{k=t-n}^t f(k)$ with n odd),*

$$|\langle (P^t)_{ij} \rangle_t - \pi_j| \leq \sqrt{\frac{d_j}{d_i}} \langle \nu_\star^t \rangle_t. \quad (10.1.1)$$

Proof: It is a fact that for a bipartite graph, $1 - \lambda_{|V|} = -1$. From Proposition 9, we then have

$$\langle (P^t)_{ij} \rangle_t = \pi_j + \sum_{k=2}^{|V|-1} \langle (1 - \lambda_k)^t \rangle_t \cdot v_i^{(k)} v_j^{(k)} \sqrt{\frac{d_j}{d_i}},$$

since successive terms involving $(1 - \lambda_{|V|})^t = (-1)^t$ cancel when averaging over an even number of t .

The rest of the proof is analogous to the proof of Theorem 10. \square

Theorem 11 (Convergence Theorem; Theorem 4.9 in [42]). *Suppose an irreducible and aperiodic Markov Chain with transition matrix $P \in \mathbb{R}^{|V| \times |V|}$ has stationary distribution $\pi \in \mathbb{R}^{|V|}$. Then there exist constants $\alpha \in (0, 1)$, $C > 0$ such that*

$$\sup_{i \in S} \|(P^t)_i \cdot \pi\|_{TV} \leq C\alpha^t,$$

where $\|\cdot\|_{TV}$ is the total variation distance of two distributions μ, ν on S , $\|\mu - \nu\|_{TV} = \max_{A \subseteq S} |\mu(A) - \nu(A)|$.

Proof: Since P is irreducible and aperiodic, there exists an r such that P^r has strictly positive entries. Let Π be the $|V| \times |V|$ matrix whose rows are the stationary distribution π . For sufficiently small $\delta > 0$ we have

$$(P^r)_{ij} \geq \delta \pi_j \quad (A.15)$$

for all $i, j = 1, 2, \dots, |V|$. Let $\theta = 1 - \delta$. Then the equation

$$P^r = (1 - \theta)\Pi + \theta Q \quad (A.16)$$

defines a stochastic matrix Q . It is a straightforward computation to check that $M\Pi = \Pi$ for any stochastic matrix M and that $\Pi M = \Pi$ for any matrix M such that $\pi M = \pi$.

Next, we use induction to demonstrate that

$$P^{rk} = (1 - \theta^k)\Pi + \theta^k Q^k \quad (A.17)$$

for $k \geq 1$. If $k = 1$ this holds by (A.16). Assuming that (A.17) holds for $k = n$,

$$P^{r(n+1)} = P^{rn} P^r = [(1 - \theta^n)\Pi + \theta^n Q^n] P^r. \quad (A.18)$$

Distributing and expanding P^r in the second term, using (A.16), gives

$$P^{r(n+1)} = (1 - \theta^n)\Pi P^r + (1 - \theta)\theta^n Q^n \Pi + \theta^{n+1} Q^{n+1}. \quad (A.19)$$

Since $\Pi P^r = \Pi$ and $Q^n \Pi = \Pi$, we have

$$P^{r(n+1)} = (1 - \theta^{n+1})\Pi + \theta^{n+1} Q^{n+1} \quad (A.20)$$

which establishes (A.17) for all $k \geq 1$. Multiplying by P^j and rearranging terms now yields

$$P^{rk+j} - \Pi = \theta^k (Q^k P^j - \Pi). \quad (A.21)$$

To complete the proof, sum the absolute values of the elements in row i on both sides and divide by 2. On the right-hand-side, the second factor is at most the largest possible

total variation distance between distributions, which is 1. Hence for any i , we have

$$\|(P^{rk+j})_i \cdot \pi\|_{TV} \leq \theta^k \quad (A.22)$$

and the theorem follows. \square

ACKNOWLEDGMENTS

DK is funded by a Studentship from the UK Engineering and Physical Sciences Research Council (EPSRC). AIAR gratefully acknowledges the financial support of the CMIH and CCIMI University of Cambridge. DH is funded by EPSRC grant EP/L016516/1 for the University of Cambridge Centre for Doctoral Training, the Cambridge Centre for Analysis.

REFERENCES

- [1] J. Neville and D. Jensen, "Iterative classification in relational data," *Proc. AAAI-2000 workshop on learning statistical models from relational data*, 2000.
- [2] D. Liben-Nowell and J. Kleinberg, "The link-prediction problem for social networks," *Journal of the American society for information science and technology*, vol. 58, no. 7, pp. 1019–1031, 2007.
- [3] L. Akoglu, H. Tong, and D. Koutra, "Graph based anomaly detection and description: a survey," *Data mining and knowledge discovery*, vol. 29, no. 3, pp. 626–688, 2015.
- [4] F. Nie, W. Zhu, and X. Li, "Unsupervised large graph embedding," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 31, no. 1, 2017.
- [5] A. I. Aviles-Rivero, N. Papadakis, R. Li, S. M. Alsaleh, R. T. Tan, and C.-B. Schonlieb, "When labelled data hurts: Deep semi-supervised classification with the graph 1-laplacian," *arXiv preprint arXiv:1906.08635*, 2019.
- [6] M. Balasubramanian, E. L. Schwartz, J. B. Tenenbaum, V. de Silva, and J. C. Langford, "The isomap algorithm and topological stability," *Science*, vol. 295, no. 5552, 2002.
- [7] W. N. Anderson Jr and T. D. Morley, "Eigenvalues of the laplacian of a graph," *Linear and multilinear algebra*, vol. 18, no. 2, pp. 141–145, 1985.
- [8] S. T. Roweis and L. K. Saul, "Nonlinear dimensionality reduction by locally linear embedding," *science*, vol. 290, no. 5500, pp. 2323–2326, 2000.
- [9] M. Ou, P. Cui, J. Pei, Z. Zhang, and W. Zhu, "Asymmetric transitivity preserving graph embedding," in *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, 2016, pp. 1105–1114.
- [10] T. Pang, F. Nie, and J. Han, "Flexible orthogonal neighborhood preserving embedding," in *IJCAI*, 2017, pp. 2592–2598.
- [11] T. M. Le and H. W. Lauw, "Probabilistic latent document network embedding," in *2014 IEEE International Conference on Data Mining*, 2014, pp. 270–279.
- [12] B. Alharbi and X. Zhang, "Learning from your network of friends: A trajectory representation learning model based on online social ties," in *2016 IEEE 16th International Conference on Data Mining (ICDM)*, 2016, pp. 781–786.
- [13] H. Xiao, M. Huang, L. Meng, and X. Zhu, "Ssp: semantic space projection for knowledge graph embedding with text descriptions," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 31, no. 1, 2017.
- [14] X. Wei, L. Xu, B. Cao, and P. S. Yu, "Cross view link prediction by learning noise-resilient representation consensus," in *Proceedings of the 26th International Conference on World Wide Web*, 2017, pp. 1611–1619.
- [15] S. Guo, Q. Wang, B. Wang, L. Wang, and L. Guo, "Semantically smooth knowledge graph embedding," in *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 2015, pp. 84–94.
- [16] S. F. Mousavi, M. Safayani, A. Mirzaei, and H. Bahonar, "Hierarchical graph embedding in vector space by graph pyramid," *Pattern Recognition*, vol. 61, pp. 245–254, 2017.
- [17] C. Yang, Z. Liu, D. Zhao, M. Sun, and E. Y. Chang, "Network representation learning with rich text information," in *IJCAI*, vol. 2015, 2015, pp. 2111–2117.

- [18] B. Perozzi, R. Al-Rfou, and S. Skiena, "Deepwalk: Online learning of social representations," in *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, 2014, pp. 701–710.
- [19] L. Van der Maaten and G. Hinton, "Visualizing data using t-sne," *Journal of machine learning research*, vol. 9, no. 11, 2008.
- [20] A. Grover and J. Leskovec, "node2vec: Scalable feature learning for networks," in *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, 2016, pp. 855–864.
- [21] Z. Yang, W. Cohen, and R. Salakhudinov, "Revisiting semi-supervised learning with graph embeddings," in *International conference on machine learning*. PMLR, 2016, pp. 40–48.
- [22] S. Cao, W. Lu, and Q. Xu, "Deep neural networks for learning graph representations," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 30, no. 1, 2016.
- [23] T. Mikolov, K. Chen, G. S. Corrado, and J. Dean, "Efficient estimation of word representations in vector space," *CoRR*, vol. abs/1301.3781, 2013.
- [24] B. Perozzi, V. Kulkarni, H. Chen, and S. Skiena, "Don't walk, skip! online learning of multi-scale network embeddings," in *Proceedings of the 2017 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2017*, 2017, pp. 258–265.
- [25] C. Tu, W. Zhang, Z. Liu, M. Sun *et al.*, "Max-margin deepwalk: Discriminative learning of network representation." in *IJCAI*, vol. 2016, 2016, pp. 3889–3895.
- [26] 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 *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*, 2018, pp. 459–467.
- [27] F. Chen, B. Wang, and C.-C. J. Kuo, "Deepwalk-assisted graph pca (dgpca) for language networks," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 2957–2961.
- [28] Z.-H. Chen, Z.-H. You, Z.-H. Guo, H.-C. Yi, G.-X. Luo, and Y.-B. Wang, "Prediction of drug–target interactions from multi-molecular network based on deep walk embedding model," *Frontiers in Bioengineering and Biotechnology*, vol. 8, p. 338, 2020.
- [29] H. Chen, B. Perozzi, Y. Hu, and S. Skiena, "Harp: Hierarchical representation learning for networks," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32, no. 1, 2018.
- [30] T. Pimentel, A. Veloso, and N. Ziviani, "Unsupervised and scalable algorithm for learning node representations," *International Conference of Learning Representation*, 2017.
- [31] Z. Yang, J. Tang, and W. Cohen, "Multi-modal bayesian embeddings for learning social knowledge graphs," *International Joint Conference on Artificial Intelligence*, 2015.
- [32] S. Pan, J. Wu, X. Zhu, C. Zhang, and Y. Wang, "Tri-party deep network representation," *International Joint Conference on Artificial Intelligence*, 2016.
- [33] J. Li, J. Zhu, and B. Zhang, "Discriminative deep random walk for network classification," in *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2016, pp. 1004–1013.
- [34] J. Tang, M. Qu, M. Wang, M. Zhang, J. Yan, and Q. Mei, "Line: Large-scale information network embedding," in *Proceedings of the 24th international conference on world wide web*, 2015, pp. 1067–1077.
- [35] S. Cao, W. Lu, and Q. Xu, "Grarep: Learning graph representations with global structural information," in *Proceedings of the 24th ACM international on conference on information and knowledge management*, 2015, pp. 891–900.
- [36] H. Chen, S. F. Sultan, Y. Tian, M. Chen, and S. Skiena, "Fast and accurate network embeddings via very sparse random projection," in *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, ser. CIKM '19. New York, NY, USA: Association for Computing Machinery, 2019, p. 399–408.
- [37] J. Qiu, C. Wang, B. Liao, R. Peng, and J. Tang, "A matrix chernoff bound for markov chains and its application to co-occurrence matrices," *Advances in Neural Information Processing Systems*, vol. 33, 2020.
- [38] Y. Zhang and M. Tang, "Consistency of random-walk based network embedding algorithms," *arXiv preprint arXiv:2101.07354*, 2021.
- [39] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean, "Distributed representations of words and phrases and their compositionality," in *Advances in neural information processing systems*, 2013, pp. 3111–3119.
- [40] J. R. Norris, *Markov Chains*, ser. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 1997, p. 1–59.
- [41] L. Lovász, "Random walks on graphs: A survey," *Combinatorics, Paul Erdős is Eighty*, vol. 2, no. 1, pp. 1–46, 1993.
- [42] D. A. Levin, Y. Peres, and E. L. Wilmer, *Markov Chains and Mixing Times*. American Mathematical Soc., 2017, vol. 107.
- [43] P. Diaconis and D. Stroock, "Geometric bounds for eigenvalues of markov chains," *The Annals of Applied Probability*, pp. 36–61, 1991.
- [44] L. Tang and H. Liu, "Relational learning via latent social dimensions," in *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, 2009, pp. 817–826.
- [45] P. Sen, G. Namata, M. Bilgic, L. Getoor, B. Galligher, and T. Eliassi-Rad, "Collective classification in network data," *AI magazine*, vol. 29, no. 3, pp. 93–93, 2008.
- [46] B. Rozemberczki, C. Allen, and R. Sarkar, "Multi-scale attributed node embedding," *arXiv preprint arXiv:1909.13021*, 2019.
- [47] A. Mislove, M. Marcon, K. P. Gummadi, P. Druschel, and B. Bhat-tacharjee, "Measurement and analysis of online social networks," in *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*, 2007, pp. 29–42.