

Learning Role-based Graph Embeddings

Nesreen K. Ahmed
Intel Labs

Ryan A. Rossi
Adobe Labs

John Boaz Lee
WPI

Xiangnan Kong
WPI

Theodore L. Willke
Intel Labs

Rong Zhou
Google

Hoda Eldardiry
PARC

Abstract

Random walks are at the heart of many existing network embedding methods. However, such algorithms have many limitations that arise from the use of random walks, *e.g.*, the features resulting from these methods are unable to transfer to new nodes and graphs as they are tied to vertex identity. In this work, we introduce the *Role2Vec* framework which uses the flexible notion of *attributed random walks*, and serves as a basis for generalizing existing methods such as DeepWalk, node2vec, and many others that leverage random walks. Our proposed framework enables these methods to be more widely applicable for both transductive and inductive learning as well as for use on graphs with attributes (if available). This is achieved by learning functions that generalize to new nodes and graphs. We show that our proposed framework is effective with an average AUC improvement of 16.55% while requiring on average 853x less space than existing methods on a variety of graphs.

1 Introduction

Learning a useful feature representation from graph data lies at the heart and success of many machine learning tasks such as node classification [Neville and Jensen, 2000], anomaly detection [Akoglu *et al.*, 2015], link prediction [Al Hasan and Zaki, 2011], among others [Koyutürk *et al.*, 2006; Ng *et al.*, 2002]. Motivated by the success of word embedding models, such as the skip-gram model [Mikolov *et al.*, 2013b; 2013a], recent works extended word embedding models to learn graph embeddings [Perozzi *et al.*, 2014; Goyal and Ferrara, 2017]. The primary goal of these works is to model the conditional probabilities that relate each input vertex to its context, where the context is a set of other vertices surrounding and/or topologically related to the input vertex. Many variants of graph embedding methods proposed *random walks* to generate the context vertices [Perozzi *et al.*, 2014; Grover and Leskovec, 2016; Ribeiro *et al.*, 2017; Cavallari *et al.*, 2017]. For instance, DeepWalk [Perozzi *et al.*, 2014] initiates random walks from each vertex to collect sequences of vertices (similar to sentences in language). Then, the skip-gram model is used to fit the embeddings by

maximizing the conditional probabilities that relate each input vertex to its surrounding context. In this case, vertex identities are used as words in the skip-gram model, and the embeddings are tied to the vertex ids.

In language, the foundational idea is that words with similar meanings will be surrounded by a similar context [Harris, 1954]. As such, in language models, the context of a word is defined as the surrounding words. However, this foundation does not directly translate to graphs. Since unlike words in languages that are universal with semantics and meaning independent of the corpus of documents, vertex ids obtained by random walks on graphs are not universal and are only meaningful within a particular graph. This key limitation has two main disadvantages. First, these embedding methods are inherently transductive, dealing essentially with isolated graphs, and unable to generalize to unseen nodes. Consequently, they are unsuitable for graph-based transfer learning tasks such as across-network classification [Kuwadekar and Neville, 2011; Getoor and Taskar, 2007], and graph similarity/comparison [Goldsmith and Davenport, 1990; Zager and Verghese, 2008]. Second, by using this traditional definition of random walks, there is no general way to integrate vertex attributes/features to the network representation.

While similar words are typically surrounded by a similar context, there is no guarantee that similar vertices are surrounded by similar context (obtained using random walks on graphs). Recent empirical analysis shows that using random walks in graph embeddings primarily capture proximity among the vertices (see [Goyal and Ferrara, 2017]), so that vertices that are close to one another in the graph are embedded together, *e.g.*, vertices that belong to the same community are embedded similarly. Although proximity among the vertices does not guarantee their similarity, the idea of a network position or a *role* [Lorrain and White, 1977; Rossi and Ahmed, 2015a; Henderson *et al.*, 2011] is more suitable to represent the similarity and structural relatedness among vertices. Roles represent vertex connectivity patterns such as hubs, star-centers, star-edge nodes, near-cliques or vertices that act as bridges to different regions of the graph. Intuitively, two vertices belong to the same role if they are structurally similar. Random walks will likely visit nearby vertices first, which makes them suitable for finding communities, rather than roles (structural similarity) (see Sec. 3 for theoretical analysis).

To overcome the above problems, we propose the *Role2Vec* framework which serves as a basis for generalizing many existing methods that use traditional random walks. *Role2Vec* utilizes the flexible notion of attributed random walks that is not tied to vertex identity and is instead based on a function $\Phi : \mathbf{x} \rightarrow w$ that maps a vertex attribute vector to a type, such that two vertices belong to the same type if they are structurally similar. The proposed framework provides a number of important advantages to any method generalized using it. First, the proposed framework is naturally inductive as the learned features generalize to new nodes and across graphs and therefore can be used for transfer learning tasks. Second, they are able to capture structural similarity (roles) better. Third, the proposed framework is inherently space-efficient since embeddings are learned for types (as opposed to vertices) and therefore requires significantly less space than existing methods. Fourth, the proposed framework naturally supports graphs with attributes (if available/given as input). Furthermore, our approach is shown to be effective with an average improvement of 16.55% in AUC while requiring on average 853x less space than existing methods on a variety of graphs from different application domains.

2 Framework

We introduce the components of the *Role2Vec* framework.

2.1 Embedding Model

We consider an (un)directed input graph $G = (V, E)$, where $N_v = |V|$ is the number of vertices in G , and $N_e = |E|$ is the number of edges in G . For any vertex $v_i \in V$, let $\Gamma(i)$ be the set of direct neighbors of v_i , and $d_i = |\Gamma(i)|$ is the vertex degree. In addition, we consider a matrix \mathbf{X} of attributes/features, where each \mathbf{x}_i is a K -vector for vertex v_i . For example, for graphs without attributes, \mathbf{x}_i could simply be an indicator vector for vertex v_i and K is equivalent to the number of vertices (*i.e.*, having $x_{ij} = 1$ if $j = i$, and $x_{ij} = 0$ otherwise) [Perozzi *et al.*, 2014; Grover and Leskovec, 2016]. For attributed graphs, \mathbf{x}_i may include observed attributes, topological features, and/or node types for heterogeneous graphs. The goal of an embedding method is to derive useful features of particular graph elements (*e.g.*, vertices, edges) by learning a model that maps each graph element to the latent D -dimension space. While the approach remains general for any graph element, this paper focuses on vertex embeddings.

To achieve this, an embedding is usually defined with three components: (1) the context function, which specifies a set of other vertices called the *context* c_i for any given vertex v_i , such that the context vertices are surrounding and/or topologically related to the given vertex. Each vertex is associated with two latent vectors, an *embedding* vector $\alpha_i \in \mathbb{R}^D$ and a *context* vector $\beta_i \in \mathbb{R}^D$. (2) the conditional distribution, which specifies the statistical distribution used to combine the embedding and context vectors. More specifically, the conditional distribution of a vertex combines its embedding and the context vectors of its surrounding vertices. (3) the model parameters (*i.e.*, embedding and context vectors) and how these are shared across the conditional distributions. Thus, an embedding method models the conditional probability that relate

each vertex to its context as follows: $\mathbf{x}_{c_i} | \mathbf{x}_i \sim \mathbb{P}$, where c_i is the set of context vertices for vertex v_i , \mathbf{x}_i is its feature/attribute vector, and \mathbb{P} is the conditional distribution.

Our goal is to model $\mathbb{P}[\mathbf{x}_{c_i} | \mathbf{x}_i] = \prod_{j \in c_i} \mathbb{P}(\mathbf{x}_j | \mathbf{x}_i)$, assuming the context vertices are conditionally independent. The most commonly used conditional distribution is the categorical distribution (see [Rudolph *et al.*, 2016] for other distributions). In this case, a softmax function parameterized with the two latent vectors (*i.e.*, embedding and context vectors) is used. Thus, for each input-context vertex pair (v_i, v_j) ,

$$\mathbb{P}(\mathbf{x}_j | \mathbf{x}_i) = \frac{e^{\alpha_i \cdot \beta_j}}{\sum_{v_k \in V} e^{\alpha_i \cdot \beta_k}} \quad (1)$$

For sparse graphs, the summation in $\sum_{v_k \in V} e^{\alpha_i \cdot \beta_k}$ contains many zero entries, and thus can be approximated by subsampling those zero entries (using negative sampling similar to language models [Mikolov *et al.*, 2013b]). Finally, the objective function of the embedding method is the sum of the logarithm of likelihood values of each vertex, *i.e.*, $\mathcal{L}(\alpha, \beta) = \sum_{i=1}^{N_v} \log \mathbb{P}[\mathbf{x}_{c_i} | \mathbf{x}_i]$.

Clearly, there is a class of possible embedding methods where each of the three components (discussed above) is considered a modeling choice with various alternatives. Recent work proposed *random walks* to sample/collect the context vertices c_i [Perozzi *et al.*, 2014; Grover and Leskovec, 2016]. Note for these random walk based embedding methods, \mathbf{x}_i is simply an indicator vector for vertex v_i (*i.e.*, no attributes).

2.2 Mapping Vertices to Vertex-Types

Given $N_v \times K$ matrix \mathbf{X} of attributes and/or structural features, the *Role2Vec* framework starts by locating sets of vertices, however large or small be the shortest distance between any two in a set, who are placed similarly with respect to all other sets of vertices. Thus, two vertices belong to the same set if they are similar in terms of attributes and/or structural features. We achieve this by learning a function that maps the N_v vertices to a set $W = \{w_1, \dots, w_M\}$ of M *vertex-types* where M is often much smaller than N_v , *i.e.*, $M \ll N_v$,

$$\Phi : \mathbf{x} \rightarrow w \quad (2)$$

Thus, Φ is a function mapping vertices to *vertex-types* based on the $N_v \times K$ attribute matrix \mathbf{X} . Clearly, the function Φ is a modeling choice, which could be learned automatically or defined manually by the user. We explore two general classes of functions for mapping vertices to their types. The first class of functions are simple functions taking the form:

$$\Phi(\mathbf{x}) = x_1 \circ x_2 \circ \dots \circ x_K \quad (3)$$

where $\mathbf{x} = [x_1 \ x_2 \ \dots \ x_K]$ is an attribute vector and \circ is a binary operator such as concatenation, sum, among others.

The second class of functions are learned by solving an objective function. This includes functions based on a low-rank factorization of the $N_v \times K$ matrix \mathbf{X} having the form $\mathbf{X} \approx f(\mathbf{U}\mathbf{V}^T)$ with factor matrices $\mathbf{U} \in \mathbb{R}^{N_v \times r}$ and $\mathbf{V} \in \mathbb{R}^{K \times r}$ where $r > 0$ is the rank and f is a linear or non-linear function. More formally,

$$\arg \min_{\mathbf{U}, \mathbf{V} \in \mathcal{C}} \left[\mathcal{D}(\mathbf{X}, f(\mathbf{U}\mathbf{V}^T)) + \mathcal{R}(\mathbf{U}, \mathbf{V}) \right] \quad (4)$$

where \mathcal{D} is the loss, \mathcal{C} is constraints (*e.g.*, non-negativity constraints $\mathbf{U} \geq 0, \mathbf{V} \geq 0$), and $\mathcal{R}(\mathbf{U}, \mathbf{V})$ is a regularization penalty. Then, we partition $\mathbf{U} \in \mathbb{R}^{N_v \times r}$ into M disjoint sets of nodes (for each of the M *vertex-types*) V_1, \dots, V_M , where

V_j is set of vertices mapped to vertex-type $w_j \in W$, by solving the k-means objective:

$$\min_{\{V_j\}_{j=1}^M} \sum_{j=1}^M \sum_{\mathbf{u}_i \in V_j} \|\mathbf{u}_i - \mathbf{c}_j\|^2, \text{ where } \mathbf{c}_j = \frac{\sum_{\mathbf{u}_i \in V_j} \mathbf{u}_i}{|V_j|} \quad (5)$$

2.3 Attributed Random Walks

Recently, random walks received much attention in learning network embeddings [Perozzi *et al.*, 2014; Grover and Leskovec, 2016], in particular to generate the context vertices (see Sec 2.1). Consider a random walk of length L and starting at a vertex v_0 of the input graph G , if at time t we are at vertex v_t , then at time $t+1$, we move to a neighbor of v_t with probability $1/d_{v_t}$. Thus, the resulting randomly chosen sequence of vertex indices ($v_t : t = 0, 1, \dots, L-1$) is a Markov chain. However, a key limitation of these methods is that the embeddings learned based on random walks are fundamentally tied to vertex ids (see Sec 2.1). By using this traditional definition of random walks, there is no general way to integrate vertex attributes and structural features to the network representation. On the other hand, vertex attributes and structural features can easily be represented by differentiating the edges according to the types of their endpoints, which leads to the definition of *attributed random walks*.

Definition 1 (Attributed walk) Let \mathbf{x}_i be a K -dimensional vector for vertex v_i . An attributed walk of length L is a sequence of adjacent vertex-types,

$$\Phi(\mathbf{x}_{v_0}), \dots, \Phi(\mathbf{x}_{v_t}), \Phi(\mathbf{x}_{v_{t+1}}), \dots, \Phi(\mathbf{x}_{v_{L-1}}) \quad (6)$$

induced by a randomly chosen sequence of indices ($v_t : t = 0, 1, \dots, L-1$) generated by a random walk of length L starting at v_0 , and a function $\Phi : \mathbf{x} \rightarrow w$ that maps an input vector \mathbf{x} to a vertex type $\Phi(\mathbf{x})$.

The induced vertex-type sequence in the above definition is called *attributed random walks* and is also a Markov chain.

The *Role2vec* framework uses vertex mapping and attributed random walks to learn the embeddings. Thus, our goal is to model the conditional probability that relate each vertex-type to the types of its context,

$$\mathbb{P}[\Phi(\mathbf{x}_{c_i}) | \Phi(\mathbf{x}_i)] = \prod_{j \in c_i} \mathbb{P}(\Phi(\mathbf{x}_j) | \Phi(\mathbf{x}_i)) \quad (7)$$

Hence, the embedding structure (*i.e.*, the embedding and context vectors) is shared among the vertices with the same vertex-type. Specifically, we learn $\alpha_j \in \mathbb{R}^D$ and $\beta_j \in \mathbb{R}^D$ for each partition V_j of vertices, which are mapped to vertex-type w_j . Note that *Role2vec* learns an embedding for an aggregated network, where detailed relations among individual vertices are aggregated to total relations among vertex-types.

2.4 Role2Vec Algorithm

The *Role2Vec* algorithm is shown in Alg. 1. Alg. 1 takes the following inputs: (1) graph G , (2) attribute matrix \mathbf{X} , (3) embedding dimension D , (4) walks per vertex R , (5) walk length L , (6) context window size ω . In Line 3, if \mathbf{X} is not available, we derive structural features using the graph structure itself. For instance, in this paper, we use small subgraphs called motifs as structural features. Counts of motif patterns were shown to be useful for a variety of network analysis

Algorithm 1 Role2Vec

```

1 procedure ROLE2VEC( $G = (V, E)$  and  $\mathbf{X}$ , embedding dimensions  $D$ ,
  walks per node  $R$ , walk length  $L$ , context (window) size  $\omega$ )
2   Initialize set of attributed walks  $S$  to  $\emptyset$ 
3   Extract (motif) features if needed and append to  $\mathbf{X}$ 
4   Transform each attribute in  $\mathbf{X}$  (e.g., using logarithmic binning)
5   Map vertices to types function  $\Phi : \mathbf{x} \rightarrow w$ 
6   Precompute transition probabilities  $\pi$ 
7    $G' = (V, E, \pi)$ 
8   parallel for  $j = 1, 2, \dots, R$  do ▷ walks per node
9     Set  $\Pi$  to be a random permutation of the nodes  $V$ 
10    for each  $v \in \Pi$  in order do
11       $S = \text{ATTRIBUTEDWALK}(G', \mathbf{X}, v, \Phi, L)$ 
12      Add the attributed walk  $S$  to  $S$ 
13    end parallel
14     $\mathbf{A} = \text{STOCHASTICGRADIENTDESCENT}(\omega, D, S)$  ▷ parallel
15  return the learned type embeddings  $\mathbf{A}$ 

```

Algorithm 2 Attributed Random Walk

```

1 procedure ATTRIBUTEDWALK( $G', \mathbf{X}$ , start node  $s$ , function  $\Phi$ ,  $L$ )
2   Initialize attributed walk  $S$  to  $[\Phi(\mathbf{x}_s)]$ 
3   Set  $i = s$  ▷ current node
4   for  $\ell = 1$  to  $L - 1$  do
5      $\Gamma_i = \text{Set of the neighbors for node } i$ 
6      $j = \text{ALIASSAMPLE}(\Gamma_i, \pi)$  ▷ select node  $j \in \Gamma_i$ 
7     Append  $\Phi(\mathbf{x}_j)$  to  $S$ 
8     Set  $i$  to be the current node  $j$ 
9   return attributed walk  $S$  of length  $L$  rooted at node  $s$ 

```

tasks and can be computed quickly and efficiently with parallel algorithms [Ahmed *et al.*, 2015; Benson *et al.*, 2016; Ahmed *et al.*, 2016a; 2016b]. Since many graph properties including motifs exhibit power law distributions, we preprocess \mathbf{X} using logarithmic binning, similar to [Henderson *et al.*, 2011] (Line 4). In Line 5, vertices are mapped to vertex-types using a function $\Phi(\mathbf{x})$ as discussed in Section 2.2. Then, we precompute the random walk transition probabilities π , which could be uniform or weighted (Line 6). Finally, Lines 8-13 initiate random walks from each vertex using the notion of attributed random walks in Alg. 2. Finally, *Role2Vec* learns the embeddings using stochastic gradient descent in Line 14.

Recall that N_v is the number of nodes, M is the number of types, and $M \ll N_v$. *Role2vec* has the following properties.

Property 2.1 *Role2vec is space-efficient with space complexity $\mathcal{O}(MD + N_v)$.*

PROOF. To store the learned embeddings of the vertex-types, *Role2vec* takes $\mathcal{O}(MD)$ space. Also, *Role2vec* takes $\mathcal{O}(N)$ space for a hash table mapping vertices to their corresponding types. Thus, the total space used by *Role2vec* is $\mathcal{O}(MD + N)$, less space compared to baselines that require $\mathcal{O}(N_v D)$. \square

As $M \rightarrow N_v$, *Role2vec* converges to the baseline random walk methods [Perozzi *et al.*, 2014; Grover and Leskovec, 2016], since each vertex is mapped to a new type that uniquely identifies it from other vertices, *i.e.*, Φ is a one-to-one function from V onto itself.

3 Properties & Analysis

We consider a random walk of length L and starting at vertex v_0 of G , if at time t we are at vertex v_t , then at time $t + 1$, we move to a neighbor of v_t with probability $1/d_{v_t}$. Clearly, the randomly chosen sequence of vertex indices ($v_t : t = 0, 1, \dots, L - 1$) is a Markov chain. We denote by P_t the distribution of v_t , where $P_i(t) = \Pr(v_t = i)$ is the probability that the random walk visits vertex i at time t . Similarly, we denote by P_{ij} the transition probability from vertex i to vertex j in one step, where $P_{ij} = \Pr(v_t = j | v_{t-1} = i)$. Thus, the Markov property implies that this Markov chain is uniquely defined by its *one-step* transition matrix $\mathbf{P} = (P_{ij})_{v_i, v_j \in V}$,

$$P_{ij} = \begin{cases} 1/d_i, & \text{if } (i, j) \in E \\ 0, & \text{otherwise} \end{cases} \quad (8)$$

Let \mathbf{P}^m be the transition matrix whose entries are the m -step transition probabilities, such that

$$P_{ij}^m = \Pr(v_{t+m} = j \mid v_t = i) \quad (9)$$

is the probability that the walk moves from vertex i to vertex j in exactly m steps. Finally, we denote by r_{ij}^t the probability that starting at vertex i , the *first transition* to vertex j occurs at time t ,

$$r_{ij}^t = \Pr(v_t = j, \forall 1 \leq m \leq t - 1, v_m \neq j \mid v_0 = i) \quad (10)$$

Properties. In a graph G , the sequence of vertices visited by a random walk of length L is represented by a directed path on the graph. Thus, it is crucial to analyze the properties and parameters of random walks that affect the embedding methods. Lemmas 1–3 analyze the constraints and bounds on vertex reachability, expected access time, and representation of vertices/edges in random walks.

Lemma 1 *If u and v are two non-adjacent vertices in a connected graph G , then there is at least one neighbor $j \in \Gamma(u)$ where $r_{uj}^t \leq r_{jv}^{t_1}$ for $t_1 < t$.*

PROOF. Let $d_u = |\Gamma(u)|$ be the degree of vertex u , and denote by $[d_u]$ the set of neighbors of u . For each neighbor $j \in [d_u]$, start a random walk at j , and let $r_{jv}^{t_1}$ be the probability that the first transition from j to v occurs at time t_1 . Now begin a random walk at u and let r_{uv}^t be the probability that the first transition from u to v occurs at time t . By conditioning on the first transition, we have

$$r_{uv}^t = \sum_{j=1}^{d_u} P_{uj} \cdot r_{jv}^{t-1} = \frac{1}{d_u} \sum_{j=1}^{d_u} r_{jv}^{t-1}$$

Set $t_1 = t - 1$, thus the probability r_{uv}^t is the mean of the probabilities of u 's neighbors, $r_{jv}^{t_1}$ for $j \in [d_u]$ and $t_1 < t$. This implies that there is at least one neighbor j where $r_{uv}^t \leq r_{jv}^{t_1}$ for $t_1 < t$, and Property 1 is proved. \square

Lemma 1 shows that the probability r_{uv}^t is upper bounded by the probability of at least one of u 's neighbors (i.e., $r_{jv}^{t_1}$). This verifies the intuition that a random walk starting at any vertex u will likely visit *nearby* vertices first before visiting *distant* vertices. And that a distant vertex v is more reachable from some neighbor j in less steps (i.e., $t_1 < t$).

Lemma 2 *If u and v are two non-adjacent vertices in a connected graph G , with h_{uv} is the expected access time from u to v , and \tilde{h}_{jv} is the average neighbor access time for u , then*

with probability less than $1/2$, a random walk starting at u takes at least $L = 2 h_{uv} > 2 \tilde{h}_{jv}$ time to reach v .

PROOF. Recall that r_{uv}^t is the probability that starting at u , the random walk first visits v at time t , then the expected access time from u to v is $\mathbb{E}[t] = h_{uv} = \sum_{t \geq 1} t \cdot r_{uv}^t$. By conditioning on the first transition, we have

$$\begin{aligned} h_{uv} &= \sum_{t \geq 1} t \cdot \sum_{j=1}^{d_u} P_{uj} \cdot r_{jv}^{t-1} \\ &= \sum_{j=1}^{d_u} P_{uj} + \sum_{t > 1} t \cdot \sum_{j=1}^{d_u} P_{uj} \cdot r_{jv}^t \\ &= 1 + \frac{1}{d_u} \sum_{j=1}^{d_u} \sum_{t > 1} t \cdot r_{jv}^t = 1 + \frac{1}{d_u} \sum_{j=1}^{d_u} h_{jv} \end{aligned}$$

where d_u is the degree of u , and h_{jv} is the expected access time for some neighbor vertex $j \in [d_u]$. Since $t \geq 0$ for any vertex in G , then by Markov's inequality, for any $L > 0$,

$$\Pr(t \geq L) \leq \frac{\mathbb{E}[t]}{L} = \frac{h_{uv}}{L}$$

Let $\tilde{h}_{jv} = \frac{1}{d_u} \sum_{j=1}^{d_u} h_{jv}$ be the average neighbor access time for vertex u . Then, with $\Pr(t \geq L) \leq 1/2$, a random walk starting at u takes at least $L = 2 h_{uv} > 2 \tilde{h}_{jv}$ to reach v . \square

Lemma 3 *Suppose we start d_u random walks of length L from any vertex $u \in V$ in G . For a given edge $e = (v, v')$, let I_e denote the total number of random walks containing e . Then, the expectation of the random variable I_e is upper bounded by L , i.e., $\mathbb{E}[I_e] \leq L$.*

PROOF. Recall that the probability of a random walk starting at u visits v at time t is $P_{uv}^t = \lambda_u \mathbf{P}^t \lambda_v^\top$, where λ_u is the indicator vector for vertex u , which equals 1 in coordinate u and 0 otherwise. Then, for a given edge $e = (v, v')$, the probability that the random walk visits v at time t and v' at time $t + 1$ is $\lambda_u \mathbf{P}^t \lambda_v^\top / d_v$ (since the transition probability from v to v' is $1/d_v$). Suppose we start d_u random walks of length L from u , let I_e denote the total number of random walks containing e , then the expectation of I_e is the sum of the probabilities that there exists a random walk visiting $e = (v, v')$ as follows

$$\begin{aligned} \mathbb{E}[I_e] &\leq \sum_{t=1}^L \sum_u d_u \lambda_u \mathbf{P}^t \lambda_v^\top / d_v \\ &= \sum_{t=1}^L \mathbf{1D} \mathbf{P}^t \lambda_v^\top / d_v = \sum_{t=1}^L \mathbf{1D} \lambda_v^\top / d_v = \sum_{t=1}^L \mathbf{1} = L \end{aligned}$$

where \mathbf{D} is the degree matrix with the i th diagonal entry is the vertex degree d_i , $\mathbf{1}$ is the unit vector with all entries equal to 1, and $(\mathbf{1D})\mathbf{P}^t = \mathbf{1D}$. \square

4 Experiments

In this section, we investigate the effectiveness of the proposed framework using a variety of graphs. Unless otherwise mentioned, all experiments use logarithmic binning¹ and the bin size δ is chosen by searching over $\delta \in$

¹Logarithmic binning assigns the first δN_v nodes with smallest attribute value to 0 (where $0 < \delta < 1$), then assigns the δ fraction of remaining unassigned nodes with smallest value to 1, and so on.

$\{0.01, 0.1, 0.5, 0.9, 0.99\}$. In these experiments, we use a simple function $\Phi(\mathbf{x})$ that represents a concatenation of the attribute values in the node attribute vector \mathbf{x} . We searched over 10 subsets of the 9 motif features of size 2-4 nodes shown in Figure 1. We evaluate the *role2vec* approach presented in Section 2.4 that leverages the attributed random walk framework (Section 2) against a number of popular methods including: *node2vec* [Grover and Leskovec, 2016], *DeepWalk* [Perozzi *et al.*, 2014], *struc2vec* [Ribeiro *et al.*, 2017], and *LINE* [Tang *et al.*, 2015]. For our approach and *node2vec*, we use the same hyperparameters ($D = 128$, $R = 10$, $L = 80$) and grid search over $p, q \in \{0.25, 0.50, 1, 2, 4\}$ as mentioned in [Grover and Leskovec, 2016]. We use logistic regression (LR) with an L2 penalty. The model is selected using 10-fold cross-validation on 10% of the labeled data. Experiments are repeated for 10 random seed initializations. All results are statistically significant with p-value < 0.01 . We use AUC to evaluate the models. Data was obtained from NetworkRepository [Rossi and Ahmed, 2015b].




Figure 1: Summary of the 9 motifs and 15 orbits with 2-4 nodes.

4.1 Comparison

This section compares the proposed approach to other embedding methods for link prediction. Given a partially observed graph G with a fraction of missing edges, the link prediction task is to predict these missing edges. We generate a labeled dataset of edges as done in [Grover and Leskovec, 2016]. Positive examples are obtained by removing 50% of edges randomly, whereas *negative examples* are generated by randomly sampling an equal number of node pairs that are not connected with an edge, *i.e.*, each node pair $(i, j) \notin E$. For each method, we learn features using the remaining graph that consists of only positive examples. Using the learned embeddings from each method, we then learn a model to predict whether a given edge in the test set exists in E or not. Notice that node embedding methods such as *DeepWalk* and *node2vec* require that each node in G appear in at least one edge in the training graph (*i.e.*, the graph remains connected), otherwise these methods are unable to derive features for such nodes. This is a significant limitation that prohibits their use in many real-world applications.

For comparison, we use the same set of binary operators [Grover and Leskovec, 2016] to construct features for the edges by combining the learned embeddings of its endpoints. The AUC results are provided in Table 1 and 2. Moreover, the AUC scores from our method are all significantly better than the other methods at $p < 0.01$. We summarize the gain/loss in predictive performance over the other methods in Figure 2. In all cases, our method achieves better predictive performance over the other methods across a wide variety of graphs with different characteristics. Overall, the mean and product binary operators give an average gain in predictive performance (over all graphs) of 11.1% and 22%, respectively.




Figure 2: AUC gain of Role2Vec (R2V) over the other methods for link prediction bootstrapped using Hadamard $\alpha_i \odot \alpha_j$.

We also investigated learning types using low-rank matrix factorization (Eq. 4) with squared loss. No regularization or constraints were used. Eq. 5 was used to partition nodes into M types. To select M a grid search is performed over $M \in \{8, 64, 512, 4096\}$. Results are provided in Table 3 and are comparable to previous results that use concatenation to derive types. Due to space, we report results for only a few graphs using the mean operator.

4.2 Space-efficient Embeddings

We now investigate the space-efficiency of the learned embeddings from the proposed framework and intermediate representation. Observe that any embedding method that implements the proposed attributed random walk framework (and intermediate representation) learns an embedding for each distinct node type $w \in W$. As described earlier in Sec. 2.4, in the worst case, an embedding is learned for each of the N_v vertices in the graph and we recover the baseline methods [Perozzi *et al.*, 2014; Grover and Leskovec, 2016] as a special case. In general, the best embedding most often lies between such extremes and therefore the embedding learned from a method implementing *Role2Vec* is often orders of magnitude smaller in size, since $M \ll N_v$.

Given an attribute vector \mathbf{x} of motif counts (Figure 1) for an arbitrary node in G , we derive embeddings using each of the following:

$$\Phi(\mathbf{x}_i = [x_2 \ x_3]), \text{ for } i = 1, \dots, N_v \quad (11)$$

$$\Phi(\mathbf{x}_i = [x_2 \ x_3 \ x_4 \ x_6 \ x_9]), \text{ for } i = 1, \dots, N_v \quad (12)$$

$$\Phi(\mathbf{x}_i = [x_2 \ x_3 \ x_4 \ x_5 \ \dots \ x_9]), \text{ for } i = 1, \dots, N_v \quad (13)$$

$$\Phi(\mathbf{x}_i = [x_1 \ x_2 \ x_3 \ x_4 \ x_5 \ \dots \ x_9]), \text{ for } i = 1, \dots, N_v \quad (14)$$

Table 1: AUC scores for various methods using $(\alpha_i + \alpha_j)/2$. Note **N2V**=node2vec, **DW**=DeepWalk and **S2V**=struc2vec.

GRAPH	R2V	R2V-DW	N2V	DW	LINE	S2V
bn-cat	0.710	0.688	0.627	0.627	0.672	0.669
bn-rat-brain	0.748	0.731	0.716	0.716	0.691	0.729
bn-rat-cerebral	0.867	0.846	0.813	0.811	0.709	0.858
ca-CSphd	0.838	0.838	0.768	0.735	0.620	0.791
eco-fweb-baydry	0.681	0.656	0.655	0.627	0.660	0.623
ia-radoslaw-email	0.867	0.847	0.756	0.745	0.769	0.857
soc-anybeat	0.961	0.960	0.854	0.848	0.850	0.883
soc-dolphins	0.656	0.597	0.580	0.498	0.551	0.590
fb-Yale4	0.793	0.793	0.742	0.728	0.763	0.758
web-EPA	0.926	0.925	0.804	0.738	0.768	0.861

Table 2: AUC scores for various methods using $\alpha_i \odot \alpha_j$. Note **N2V**=node2vec, **DW**=DeepWalk and **S2V**=struc2vec.

GRAPH	R2V	R2V-DW	N2V	DW	LINE	S2V
bn-cat	0.694	0.681	0.621	0.621	0.494	0.662
bn-rat-brain	0.775	0.775	0.715	0.715	0.562	0.736
bn-rat-cerebral	0.867	0.838	0.796	0.793	0.498	0.834
ca-CSphd	0.758	0.738	0.678	0.660	0.533	0.699
eco-fweb-baydry	0.684	0.644	0.673	0.631	0.516	0.599
ia-radoslaw-email	0.852	0.821	0.746	0.731	0.471	0.843
soc-anybeat	0.945	0.945	0.730	0.728	0.618	0.798
soc-dolphins	0.787	0.787	0.593	0.508	0.549	0.553
fb-Yale4	0.940	0.906	0.912	0.904	0.784	0.905
web-EPA	0.907	0.885	0.808	0.797	0.650	0.841

Table 3: AUC scores comparing types derived using concatenation vs. factorization.

	bn-cat	bn-rat-brain	ia-rado-email	web-EPA
R2V	0.710	0.748	0.867	0.926
R2V-Fac.	0.707	0.761	0.848	0.905

where $\Phi(\cdot)$ is a function that maps \mathbf{x}_i to a type $w \in W$. In these experiments, we use logarithmic binning (applied to each N_v -dimensional motif feature) with $\delta = 0.5$ and use Φ defined as the concatenation of the logarithmically binned attribute values. Embeddings are learned using the different subsets of attributes in Eq. (11)-(14). For instance, Eq. (11) indicates that vertex types are derived using the (logarithmic binned) number of 2-stars x_2 and triangles x_3 incident to the given vertex (Figure 1). We measure the space (in bytes) required to store the embedding learned by each method. In Figure 3, we summarize the reduction in space from our approach compared to the other methods. In all cases, the embeddings learned from our approach require significantly less space and thus more space-efficient. Specifically, the embeddings from our approach require on average 853 times less space than the best method averaged across all graphs.

5 Related Work

Random walks are widely used in graph embedding methods [Perozzi *et al.*, 2014; Grover and Leskovec, 2016; Ribeiro *et al.*, 2017; Cavallari *et al.*, 2017]. Recent embedding methods for graphs have largely been based on the popular skip-gram model [Mikolov *et al.*, 2013a; Cheng *et al.*, 2006] originally introduced for learning vector representations of words in text. In particular, DeepWalk [Perozzi *et al.*, 2014] used this approach to embed the nodes such that the co-occurrence frequencies of pairs in short random walks are preserved. Recently, node2vec [Grover and Leskovec, 2016] introduced hyperparameters to DeepWalk that tune the depth and breadth of the random walks. These approaches are becoming increasingly popular and have been shown to outperform a number of existing methods. These methods (and many others) are all based on simple random walks and thus are well-suited for generalization using the *attributed random walk framework*. While most network representation learning methods use only the graph [Perozzi *et al.*, 2014; Tang *et al.*, 2015; Cao *et al.*, 2015; Grover and Leskovec, 2016], our framework exploits both the graph *and* structural features (*e.g.*, motifs). While most work has focused on transductive (within-




Figure 3: Evaluating the space savings of Role2Vec. Space savings is defined as the reduction in embedding size (bytes) relative to the smallest embedding size among the baseline methods. Let σ denote the size (bytes) of a Role2Vec embedding and σ_b be the smallest embedding size from the baseline methods, then space savings is defined as $1 - \frac{\sigma}{\sigma_b}$. Hence, larger values indicate a larger space savings (reduction in size). Notably, the Role2Vec variants require significantly less space than existing methods as shown above. Note $[x_2 x_3], \dots, [x_1 \dots x_9]$ is the attribute sets in Eq. (11)-(14) used as input to Φ for mapping nodes to types.

network) learning, there has been some recent work on graph-based inductive approaches. Yang *et al.* [2016] proposed an inductive approach called Planetoid. However, Planetoid is an embedding-based approach for semi-supervised learning and does not use any structural features. Recently, Hamilton *et al.* [2017] proposed a similar approach that also aggregates features from node neighborhoods. However, these approaches are not based on random-walks. Heterogeneous networks [Shi *et al.*, 2014] have also been recently considered [Chang *et al.*, 2015; Dong *et al.*, 2017] as well as attributed networks [Huang *et al.*, 2017b; 2017a]. Huang *et al.* [2017b] proposed an approach for attributed networks with labels whereas Yang *et al.* [2015] used text features to learn node representations. Liang *et al.* [2017] proposed a semi-supervised approach for networks with outliers. Bojchevski *et al.* [2017] proposed an unsupervised rank-based approach. Coley *et al.* [2017] introduced a convolutional approach for attributed molecular graphs that learns graph embeddings as opposed to node embeddings. However, most of these approaches are neither inductive nor space-efficient.

6 Conclusion

This work proposed a flexible framework based on the notion of attributed random walks. The framework serves as a basis for generalizing existing techniques (that are based on random walks) for use with attributed graphs, unseen nodes, graph-based transfer learning tasks, and allowing significantly larger graphs due to the inherent space-efficiency of the approach. Instead of learning individual embeddings for each node, our approach learns embeddings for each type based on functions that map feature vectors to types. This allows for both inductive and transductive learning. Furthermore, the framework was shown to have the following desired properties: space-efficient, accurate, inductive, and able

to support graphs with attributes (if available).

References

- [Ahmed *et al.*, 2015] Nesreen K. Ahmed, Jennifer Neville, Ryan A. Rossi, and Nick Duffield. Efficient graphlet counting for large networks. In *ICDM*, page 10, 2015.
- [Ahmed *et al.*, 2016a] Nesreen K. Ahmed, Jennifer Neville, Ryan A. Rossi, Nick Duffield, and Theodore L. Willke. Graphlet decomposition: Framework, algorithms, and applications. *KAIS*, pages 1–32, 2016.
- [Ahmed *et al.*, 2016b] Nesreen K. Ahmed, Theodore L. Willke, and Ryan A. Rossi. Estimation of local subgraph counts. In *IEEE BigData*, pages 586–595, 2016.
- [Akoglu *et al.*, 2015] Leman Akoglu, Hanghang Tong, and Danai Koutra. Graph based anomaly detection and description: a survey. *DMKD*, 29(3):626–688, 2015.
- [Al Hasan and Zaki, 2011] Mohammad Al Hasan and Mohammed J Zaki. A survey of link prediction in social networks. In *Social Network Data Analytics*, pages 243–275. 2011.
- [Benson *et al.*, 2016] Austin R Benson, David F Gleich, and Jure Leskovec. Higher-order organization of complex networks. *Science*, 353(6295):163–166, 2016.
- [Bojchevski and Günnemann, 2017] Aleksandar Bojchevski and Stephan Günnemann. Deep gaussian embedding of attributed graphs: Unsupervised inductive learning via ranking. *arXiv:1707.03815*, 2017.
- [Cao *et al.*, 2015] Shaosheng Cao, Wei Lu, and Qiongkai Xu. Grarep: Learning graph representations with global structural information. In *CIKM*, pages 891–900. ACM, 2015.
- [Cavallari *et al.*, 2017] Sandro Cavallari, Vincent W Zheng, Hongyun Cai, Kevin Chen-Chuan Chang, and Erik Cambria. Learning community embedding with community detection and node embedding on graphs. In *CIKM*, pages 377–386, 2017.
- [Chang *et al.*, 2015] Shiyu Chang, Wei Han, Jiliang Tang, Guo-Jun Qi, Charu C Aggarwal, and Thomas S Huang. Heterogeneous network embedding via deep architectures. In *SIGKDD*, pages 119–128, 2015.
- [Cheng *et al.*, 2006] Winnie Cheng, Chris Greaves, and Martin Warren. From n-gram to skipgram to concgram. *Int. J. of Corp. Linguistics*, 11(4):411–433, 2006.
- [Coley *et al.*, 2017] Connor W Coley, Regina Barzilay, William H Green, Tommi S Jaakkola, and Klavs F Jensen. Convolutional embedding of attributed molecular graphs for physical property prediction. *J. Chem. Info. & Mod.*, 2017.
- [Dong *et al.*, 2017] Yuxiao Dong, Nitesh V Chawla, and Ananthram Swami. metapath2vec: Scalable representation learning for heterogeneous networks. In *SIGKDD*, pages 135–144, 2017.
- [Getoor and Taskar, 2007] L. Getoor and B. Taskar, editors. *Intro. to SRL*. MIT Press, 2007.
- [Goldsmith and Davenport, 1990] Timothy E Goldsmith and Daniel M Davenport. Assessing structural similarity of graphs. 1990.
- [Goyal and Ferrara, 2017] Palash Goyal and Emilio Ferrara. Graph embedding techniques, applications, and performance: A survey. *arXiv preprint arXiv:1705.02801*, 2017.
- [Grover and Leskovec, 2016] Aditya Grover and Jure Leskovec. node2vec: Scalable feature learning for networks. In *SIGKDD*, pages 855–864, 2016.
- [Hamilton *et al.*, 2017] William L Hamilton, Rex Ying, and Jure Leskovec. Inductive representation learning on large graphs. *arXiv:1706.02216*, 2017.
- [Harris, 1954] Zellig S Harris. Distributional structure. *Word*, 10(2-3):146–162, 1954.
- [Henderson *et al.*, 2011] Keith Henderson, Brian Gallagher, Lei Li, Leman Akoglu, Tina Eliassi-Rad, Hanghang Tong, and Christos Faloutsos. It’s who you know: graph mining using recursive structural features. In *SIGKDD*, pages 663–671. ACM, 2011.
- [Huang *et al.*, 2017a] Xiao Huang, Jundong Li, and Xia Hu. Accelerated attributed network embedding. In *SDM*, 2017.
- [Huang *et al.*, 2017b] Xiao Huang, Jundong Li, and Xia Hu. Label informed attributed network embedding. In *WSDM*, 2017.
- [Koyutürk *et al.*, 2006] Mehmet Koyutürk, Yohan Kim, Umüt Topkara, Shankar Subramaniam, Wojciech Szpankowski, and Ananth Grama. Pairwise alignment of protein interaction networks. *JCB*, 13(2):182–199, 2006.
- [Kuwadekar and Neville, 2011] Ankit Kuwadekar and Jennifer Neville. Relational active learning for joint collective classification models. In *ICML*, pages 385–392, 2011.
- [Liang *et al.*, 2017] Jiongqian Liang, Peter Jacobs, and Srinivasan Parthasarathy. Seano: Semi-supervised embedding in attributed networks with outliers. *arXiv:1703.08100*, 2017.
- [Lorrain and White, 1977] Francois Lorrain and Harrison C White. Structural equivalence of individuals in social networks. In *Social Networks*, pages 67–98. Elsevier, 1977.
- [Mikolov *et al.*, 2013a] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. In *ICLR Workshop*, page 10, 2013.
- [Mikolov *et al.*, 2013b] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119, 2013.
- [Neville and Jensen, 2000] Jennifer Neville and David Jensen. Iterative classification in relational data. In *AAAI SRL Workshop*, pages 13–20, 2000.
- [Ng *et al.*, 2002] Andrew Y Ng, Michael I Jordan, and Yair Weiss. On spectral clustering: Analysis and an algorithm. In *NIPS*, pages 849–856, 2002.
- [Perozzi *et al.*, 2014] Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. Deepwalk: Online learning of social representations. In *SIGKDD*, pages 701–710, 2014.
- [Ribeiro *et al.*, 2017] Leonardo F.R. Ribeiro, Pedro H.P. Saverese, and Daniel R. Figueiredo. Struc2vec: Learning node representations from structural identity. In *SIGKDD*, pages 385–394, 2017.
- [Rossi and Ahmed, 2015a] R.A. Rossi and N.K. Ahmed. Role discovery in networks. *TKDE*, 27(4):1112–1131, 2015.
- [Rossi and Ahmed, 2015b] Ryan A. Rossi and Nesreen K. Ahmed. The network data repository with interactive graph analytics and visualization. In *AAAI*, pages 4292–4293, 2015.
- [Rudolph *et al.*, 2016] Maja Rudolph, Francisco Ruiz, Stephan Mandt, and David Blei. Exponential family embeddings. In *NIPS*, pages 478–486, 2016.
- [Shi *et al.*, 2014] Chuan Shi, Xiangnan Kong, Yue Huang, S Yu Philip, and Bin Wu. HeteSim: A General Framework for Relevance Measure in Heterogeneous Networks. *TKDE*, 26(10):2479–2492, 2014.

- [Tang *et al.*, 2015] Jian Tang, Meng Qu, Mingzhe Wang, Ming Zhang, Jun Yan, and Qiaozhu Mei. LINE: Large-scale Information Network Embedding. In *WWW*, pages 1067–1077, 2015.
- [Yang *et al.*, 2015] Cheng Yang, Zhiyuan Liu, Deli Zhao, Maosong Sun, and Edward Y Chang. Network representation learning with rich text information. In *IJCAI*, pages 2111–2117, 2015.
- [Yang *et al.*, 2016] Zhilin Yang, William W Cohen, and Ruslan Salakhutdinov. Revisiting semi-supervised learning with graph embeddings. *arXiv:1603.08861*, 2016.
- [Zager and Verghese, 2008] Laura A Zager and George C Verghese. Graph similarity scoring and matching. *Applied mathematics letters*, 21(1):86–94, 2008.