Learning Edge Representations via Low-Rank Asymmetric Projections

Sami Abu-El-Haija
Google Research
Mountain View, California
haija@google.com

Bryan Perozzi
Google Research
New York City, New York
bperozzi@acm.org

Rami Al-Rfou
Google Research
Mountain View, California
rmyeid@google.com

ABSTRACT

We propose a new method for embedding graphs while preserving directed edge information. Learning such continuous-space vector representations (or embeddings) of nodes in a graph is an important first step for using network information (from social networks, user-item graphs, knowledge bases, etc.) in many machine learning tasks.

Unlike previous work, we (1) explicitly model an edge as a function of node embeddings, and we (2) propose a novel objective, the *graph likelihood*, which contrasts information from sampled random walks with non-existent edges. Individually, both of these contributions improve the learned representations, especially when there are memory constraints on the total size of the embeddings. When combined, our contributions enable us to significantly improve the state-of-the-art by learning more concise representations that better preserve the graph structure.

We evaluate our method on a variety of link-prediction task including social networks, collaboration networks, and protein interactions, showing that our proposed method learn representations with error reductions of up to 76% and 55%, on directed and undirected graphs. In addition, we show that the representations learned by our method are quite space efficient, producing embeddings which have higher structure-preserving accuracy but are 10 times smaller.

KEYWORDS

Graph, Edge Learning, Embedding, Random Walk, Link Prediction, Representation Learning

1 INTRODUCTION

Recent advancements in learning embedding vectors for words have resulted in a proliferation of methods which learn continuous space representations of graphs (e.g. DeepWalk [22]). These approaches process a graph and encode each node as a (real-valued) embedding vector, enabling easy integration with existing machine learning algorithms.

Such embedding methods learn a vector space that highly preserves the graph structure. Two nodes would have large similarity

Figure 1: Depiction of our method. On the left: a graph, showing a random walk in dotted-red, where nodes u, v are “close” in the walk (i.e. within a configurable context window parameter). We access the trainable embeddings Y_u and Y_v for the nodes and feed them as input to Deep Neural Network (DNN) f . The DNN outputs manifold coordinates $f(Y_u)$ and $f(Y_v)$ for nodes u and v , respectively. A low-rank asymmetric projection transforms $f(Y_u)$ and $f(Y_v)$ to their source and destination representations, which are used by g to represent an edge.

in the embedding space (or small distance) if they are *strongly* connected in the original (discrete) graph. Edges can be weighted or unweighted. Traditional *eigen* methods [3, 12, 26] learn embeddings that minimize the euclidean distance of connected nodes, which can be solved (with orthonormal constraints) by eigen-decomposition of the symmetric graph Laplacian. Recent *random-walk* embedding methods [11, 22] learn representations which encode the random walk transition matrix. These methods embed two nodes close if they co-occur frequently in short random walks. In general, *random-walk* methods outperform “eigen” methods on producing vector representations that preserve the graph structure.

However, recent random-walk embedding methods have two shortcomings. First, these methods do not explicitly model edges. This *node-centric* assumption represents an edge (u, v) identically to reverse counterpart (v, u) , and is unable to capture asymmetric relationships. Second, to preserve the graph structure they embed nodes into a relatively high-dimensional space, sometimes producing an embedding dictionary larger than the sparse adjacency matrix.

In this work we propose to address these limitations by explicitly modeling edges in the network as a function of the nodes. Specifically, we model edges by (i) using a Deep Neural Network (DNN) to map nodes onto a low-dimensional manifold, (ii) defining an

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

CIKM'17, November 6–10, 2017, Singapore, Singapore

© 2017 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-4918-5/17/11.

<https://doi.org/10.1145/3132847.3132959>

edge function between two nodes as a projection in the manifold coordinates, and (iii) jointly-optimizing the edge function and the manifold by maximizing a new objective we propose, the *graph likelihood*, which we define as a product of the edge function over all node pairs.

More formally, we learn an embedding vector $Y_u \in \mathbb{R}^D$ for every graph node u , a manifold-mapping Deep Neural Network (DNN) $f : \mathbb{R}^D \rightarrow \mathbb{R}^d$ that is shared across all nodes, and an asymmetric edge function $g : (\mathbb{R}^d \times \mathbb{R}^d) \rightarrow \mathbb{R}$ to represent edges in the graph. Our entire model $g(u, v) = f(Y_u)^T \times M \times f(Y_v)$ is end-to-end differentiable. M is low-rank, as $M = L \times R$, where both $L \in \mathbb{R}^{d \times b}$ and $R \in \mathbb{R}^{b \times d}$ project the node manifold coordinates to smaller space $\mathbb{R}^b$. Since b is much smaller than D , we are able to reduce the final node embedding significantly. Figure 1 shows a depiction of our architecture. Our desired likelihood is quadratic but we estimate it with a tractable linear objective using negative sampling, similar to [19].

We find that explicitly modeling edges can drastically reduce the representation dimensionality, for both directed and undirected graphs, especially when coupled with a Deep Neural Network. Further, modeling asymmetry by representing edge (u, v) differently than (v, u) gives an additional performance boost when preserving the structure of directed graphs. We perform an extrinsic evaluation of our method, by comparing it to the state-of-the-art on link-prediction tasks over a variety of graphs from social networks, biology, and e-commerce. We show that we can consistently learn orders of magnitude smaller embedding dimensions, while improving ROC-AUC metrics. For example, we reduce the error on directed graphs by up to $\approx 70\%$ and undirected graphs by up to $\approx 50\%$ when using same-sized representations. However, when our model is restricted to representations which are *8 times smaller* than the baselines, we reduce the error in some cases by up to 66% on directed graphs and 16% on undirected graphs. We perform intrinsic evaluations, by training and rendering two-dimensional embedding spaces for two datasets, which we use to gain intuition on placement choices made by our model.

To summarize, our contributions are as follows.

- (1) We propose to explicitly model a directed edge function, which we define as low-rank affine projections on a manifold that is produced by a Deep Neural Network (i.e. “deep embeddings”).
- (2) We propose a new objective function, the graph likelihood.
- (3) These aspects significantly improve the state-of-the-art on learning continuous graph representations, especially on directed graphs, while producing significantly smaller representation spaces, as evaluated on five graph datasets.

2 EDGE REPRESENTATIONS

It is common to embed a graph by learning one continuous D -dimensional vector $Y_u \in \mathbb{R}^D$ for every graph node $u \in V$, where relationships between nodes u and v are captured in a very coarse way, through the use of a distance measure (e.g. $\text{dist}(Y_u, Y_v)$). This *node-centric* modeling assumes that all relationships in the graph are symmetric – a limiting assumption which fails to capture any directed relationships.

We seek to model the asymmetry which occurs in many real world graphs. Specifically, given two nodes, u and v , we desire that their distances are allowed to differ ($\text{dist}(Y_u, Y_v) \neq \text{dist}(Y_v, Y_u)$) to reflect ordering in directed relationships, such as *follower* and *followee* on Twitter. In addition, the asymmetry can also model degree variance in undirected graphs. Consider a popular node m , then the optimization could make $\text{dist}(Y_u, Y_m)$ small for all u but not necessarily $\text{dist}(Y_m, Y_u)$.

Even though it is possible to learn one representation $Y_{(u,v)}$ for all node pairs (u, v) , this direct modeling is prohibitive in practice and requires an upper-bound space of $O(|V|^2)$. Instead, we propose to learn a trainable edge function defined over node embedding coordinates. Specifically, we learn asymmetric transformations of the nodes, which generates for a node u , two representations: one when it is the source of a directed edge, $\hat{Y}_u^{\text{source}}$ and one when it is a destination, $\hat{Y}_u^{\text{dest}}$. These representations share a neural network f . These representations can be combined for any pair of nodes to model the strength of their directed relationships. That is, for nodes u and v , we represent (u, v) and (v, u) as $\text{dist}(\hat{Y}_u^{\text{source}}, \hat{Y}_v^{\text{dest}})$ and $\text{dist}(\hat{Y}_v^{\text{source}}, \hat{Y}_u^{\text{dest}})$, respectively.

3 PRELIMINARIES

3.1 Link Prediction

Link prediction is a problem of inferring missing edges in a graph. We use it to evaluate the generalization ability of our embedding spaces, as we aim to preserve the graph structure. The common setup [11] is to “hold out” test edges $E_{\text{test}} \subset E$ and train on the remaining $E_{\text{train}} = E - E_{\text{test}}$. Structure-preserving representations should retrieve the held-out E_{test} with high accuracy.

3.2 Graph Embedding

Graph embedding approaches learn a D -dimensional embedding dictionary $\mathbf{Y} \in \mathbb{R}^{|V| \times D}$, containing continuous real-valued vector $Y_u \in \mathbb{R}^D$ for every graph node $u \in V$. Earlier approaches in computing embeddings include Eigenmaps [3], which embeds Y_u and Y_v to be close if they are connected (i.e. $(u, v) \in E$ or similarly $A_{uv} = 1$). Formally, Eigenmaps learns embeddings by minimizing an objective:

$$\min_{\mathbf{Y}} \sum_{(u,v) \in E} A_{uv} \|Y_u - Y_v\|_2^2 \quad \text{s.t.} \quad \mathbf{Y}^T \mathbf{D} \mathbf{Y} = \mathbf{I}, \quad (1)$$

where the weight of edge (u, v) is stored in the adjacency matrix at A_{uv} and $\mathbf{D}$ is a diagonal weight matrix with $D_{vv} = \sum_u A_{uv}$. This optimization yields an embedding space where Y_u and Y_v are near if A_{vu} is large (or non-zero). Equation (1) also appears in equivalent forms in [12, 26]. Furthermore, Bregman Iterations has been proposed to optimize an L1-formulation of the above objective function [31].

3.3 Word2vec

Word2vec [19] processes a big text corpus (e.g. Wikipedia) and learns one embedding vector for every unique word. If two words w_1 and w_2 are frequently “close” (e.g. in same sentence), then the dot product of embeddings $Y_{w_1}^T Y_{w_2}$ is maximized. In particular, every time two words w_1 and w_2 are within C words away, where integer C is the “context window size” hyperparameter, then a gradient

Algorithm 1 Extract Random Walks

Input: $G = (V, E)$, n (walks per node), τ (walk length).
Output: walks.
 $\pi = \text{makeTransitionPr}(E)$
walks = []
for $u \in V$ **do**
 for ($i = 0$; $i < n$; $i++$) **do**
 walk = [u];
 for ($j = 0$; $j < \tau$; $j++$) **do**
 curr = walk[-1]
 $v = \text{sampleNext}(\text{curr}, \pi)$
 walk.append(v)
 end for
 walks.append(walk)
 end for
end for

step increments this likelihood:

$$\frac{\exp(Y_{w_1}^T Y_{w_2})}{\sum_j \exp(Y_{w_j}^T Y_{w_2})}. \quad (2)$$

We refer the reader to [19] for further information¹

3.4 Graph Embedding with Random Walks

Rather than operating directly on the adjacency matrix A , another embedding strategy has been recently proposed by Perozzi, Al-Rfou, and Skiena [22]. Their method, DeepWalk, introduced a new class of *Random Walk* methods, which extend a node's direct neighbors to include nodes that are within small number of hops. These approaches sample many random walks from the graph. If nodes u and v are frequently close in the random walks, then the model learns a representation such that the inner product of $\langle Y_u, Y_v \rangle$ is a large positive value. Algorithm 1 extracts random walks. It begins by computing a probability transition matrix π , where $\pi_{u \rightarrow v}$ indicates the probability of a random walker visiting node v conditioned on current node being u .

This model has been extended by node2vec [11] to use a second-order probability transition function containing $\pi_{t \rightarrow u \rightarrow v}$, where the probability of a random walker visiting node v is conditioned on current node u and previous node t . Node2vec's random walk use hyper-parameters p and q , which effectively yield a graph traversal algorithm that's like an interpolation between Depth-First Search (DFS) and Breadth-First Search (BFS). We refer the reader to [11] for further details. We adopt this method for generating random walks in our work.

After sampling random walks, DeepWalk and node2vec treat each walk ($u_1 \rightarrow u_2 \rightarrow \dots \rightarrow u_\tau$) as a sequence, and then apply the skip-gram model to compute embeddings per word (i.e. node). The objective that they minimize is:

$$\min_Y \left[\log Z - \sum_{u \in V, v \in V} D_{uv} (Y_u^T Y_v) \right], \quad (3)$$

¹ The denominator, rather than summing over all words, is approximated by hierarchical softmax. In addition, their original formulation learns two vectors per word, one when used as "input" and another used when "output".

Figure 2: Depiction of f_θ , where $FC_i(x) = W_i x + b_i$ is a fully-connected layer with weight matrix W_i and bias vector b_i , bn is BatchNorm [14] and $\text{relu}(x) = \max(0, x)$ is an element-wise activation function.

Where D_{uv} is the number of times nodes u and v appear close to each other (i.e. within the context size) in all random walks. We extend these random walk methods in three important ways: First, rather than using word2vec's objective (Equation 3), we propose an novel alternative objective, the **graph likelihood** (see Section 4.2). Second, we explicitly represent an edge function as a function of nodes which we jointly train (see Section 4). Third, we define the "context" for directed graph differently than undirected ones. Specifically, a random walk $u_1 \rightarrow u_2 \rightarrow u_3 \rightarrow u_4 \rightarrow u_5$ would produce $\{u_1, \dots, u_5\}$ as a context of u_3 if graph is undirected and would produce $\{u_4, u_5\}$ as context if the graph is directed.

4 OUR METHOD

We explain the details of our model and how we train it. The source-code is made available online².

4.1 Model

Given an (un)directed graph $G = (V, E)$, we learn an embedding vector $Y_u \in \mathbb{R}^D$ for every node $u \in V$. In addition, we learn a Deep Neural Network (DNN) $f_\theta : \mathbb{R}^D \rightarrow \mathbb{R}^d$ that maps a node onto a low-dimensional manifold. f_θ is depicted in Figure 2, and is defined as:

$$f_\theta : Y_u \rightarrow FC_{\{W_1, b_1\}} \rightarrow \text{BatchNorm} \rightarrow \text{relu} \rightarrow FC_{\{W_2, b_2\}} \rightarrow \text{BatchNorm} \rightarrow f_\theta(Y_u),$$

where $FC_{\{W, b\}}$ is a fully-connected layer with weight matrix W and bias vector b , BatchNorm is described in [14], $\text{relu}(x) = \max(0, x)$ is an element-wise activation function, and $\theta = \{W_1, b_1, W_2, b_2, \dots\}$.

We define a general class of edge functions $g(u, v) \in \mathbb{R}$ where symmetry is not imposed, yielding $g(u, v) \neq g(v, u)$. Consider a low-rank affine projection in the manifold space:

$$g(u, v) = f(Y_u)^T \times M \times f(Y_v), \quad (4)$$

²Code available at http://sami.haija.org/graph/deep_embedding.html

where low-rank projection matrix $M = L \times R$ with $L \in \mathbb{R}^{d \times b}$ and $R \in \mathbb{R}^{b \times d}$. We refer to b as the bottleneck dimension and we experiment with $b < d < D$. We can factor $g(u, v)$ into an inner product $\langle L^T f(Y_u), Rf(Y_v) \rangle$. We refer to $L^T f(Y_u) \in \mathbb{R}^b$ and $Rf(Y_v) \in \mathbb{R}^b$, respectively, as the left- and right-asymmetric embeddings.

We note Equation (4) can be extended to use a combination of multiple low-rank affine projections, as:

$$g^{(2)}(u, v) = \langle w_g^{(2)}, [\text{relu}(g_1(u, v)), \dots, \text{relu}(g_h(u, v))] \rangle, \quad (5)$$

where $w_g^{(2)} \in \mathbb{R}^h$ a parameter vector of the output layer, h is the number of projections, and each projection g_i has its own $L_i \in \mathbb{R}^{d \times b}$ and $R_i \in \mathbb{R}^{b \times d}$. Even though there total size of parameters for Y, f, g may be large, we can have a low memory footprint during inference if we precompute $L_i^T f(Y_u)$ and $R_i f(Y_v)$ for every $u \in V$.

4.2 Graph Likelihood

We introduce our proposed objective, step-by-step. We start with intuitions from the Maximum Likelihood Estimate of Logistic Regression. Given a training graph $G = (V, E_{\text{train}})$, one can define a probability measure as a product of an edge estimate Q on all node pairs:

$$\Pr(G) = \prod_{(u, v) \in E_{\text{train}}} Q(u, v) \prod_{(u, v) \notin E_{\text{train}}} 1 - Q(u, v), \quad (6)$$

where $Q : V \times V \rightarrow [0, 1]$ is a trainable edge estimator. If Q is a perfect estimator, then it should output 1 on all $(u, v) \in E$ and should output 0 on all $(u, v) \notin E$, which makes $\Pr(x) = 1$ iff $x = G$. An equivalent form of equation 6 is:

$$\prod_{\substack{u \in V \\ v \in V}} Q(u, v)^{\mathbb{1}[(u, v) \in E_{\text{train}}]} (1 - Q(u, v))^{\mathbb{1}[(u, v) \notin E_{\text{train}}]} \quad (7)$$

where indicator function $\mathbb{1}[x] = 1$ if predicate x is true and is 0 otherwise. Note that two product terms are mutually exclusive, as one of the powers $\mathbb{1}[\cdot]$ will evaluate to 1 and the other to 0.

Recent work shows that extending the neighbor-set of nodes beyond their direct connections via random walks, can improve generalization of prediction tasks such as link-prediction and node classification [11, 21, 22]. Following this motivation, we propose to replace the binary edge presence $\mathbb{1}[(u, v) \in E_{\text{train}}]$ in equation (7) by simulated random walk statistics, and formulate our proposed quadratic objective, the **graph likelihood** as:

$$\Pr(G) \propto \prod_{\substack{u \in V \\ v \in V}} \sigma(g(u, v))^{\mathcal{D}_{uv}} (1 - \sigma(g(u, v)))^{\mathbb{1}[(u, v) \notin E_{\text{train}}]} \quad (8)$$

where $\sigma(x) = 1/(1 + \exp(-x))$ is the standard logistic, the edge function g is described in section 4.1, and $\mathcal{D}_{uv}$ is the unnormalized frequency that nodes u and v appear within the configured context window, in simulated random walks [22]. Note that our likelihood in equation (8) is not standard, especially that the two terms are not exclusive, 0 and $(u, v) \notin E_{\text{train}}$ to be simultaneously true. It follows that the expression under the product-operator is $\in [0, 1]$ since the logistic $\sigma : \mathbb{R} \rightarrow [0, 1]$. We use *proportional to* ($\propto$) instead of equality since the normalizing constant $\int_G \Pr(G)$ is only a scaling factor and should not change the arg max of the likelihood. Our experiments show that the likelihood yields a powerful representation

for preserving the graph structure, as evaluated on link-prediction tasks, out-performing models trained with a skipgram objective (such as node2vec [11]). We show in our Experiments (Section 5) that even when our model is identical to node2vec (i.e. shallow and symmetric), training with our proposed objective produces embeddings that better preserve the graph structure, especially when using low embedding dimensions.

Although a naïve optimization of equation (8) is quadratic, $\mathcal{D}$ is sparse with $O(|V|)$ non-zero entries, making it possible to compute the first term in equation (8) in linear time. Further, we use negative sampling (Section 4.4) to estimate the product over $(u, v) \notin E$, which is important in many real applications where graphs are large and most edges are negative, having $|\{(u, v) : u, v \in V \text{ and } (u, v) \notin E\}| \approx O(|V|)$.

4.3 Training Data Generation

Our training algorithm requires positive and negative pairs of nodes as its input. Here we briefly describe their generation.

4.3.1 Positives. Given a graph $G = (V, E)$, we take a partition $E_{\text{train}} \subset E$ and extract random walks from E_{train} using Algorithm 1. Starting from every node u_1 , we simulate n random walks, each of length τ , like:

$$u_1 \rightarrow u_2 \rightarrow u_3 \rightarrow \dots \rightarrow u_\tau.$$

Then, for every walk, we extract all node pairs within the context window, similar to [19].

$$(u_i, u_j) \quad \forall j \in \mathbb{Z}, i - w_l \leq j \leq i + w_r, j \neq i, \quad (9)$$

where w_l and w_r are the context window left and right offsets. For example, for an undirected graph, if we use a context window of size 5, then $w_l = w_r = 2$ and therefore $j \in \{-2, -1, 1, 2\}$. Extracting positive pairs yields a list $\mathcal{D}$ that contain duplicates and we overload this notation by defining $\mathcal{D}_{uv}$ as the frequency of (u, v) in list $\mathcal{D}$. It is trivial to show that the number of pairs is linear in $|V|$:

$$|\mathcal{D}| = n\tau O((w_l + w_r)(\tau - (w_l + w_r + 1))|V|) = O(|V|) \quad (10)$$

4.3.2 Negatives. We fix a set of negatives for every node. Before training, for every node u , we create its negative set $\bar{u}$ as:

$$\bar{u} = \{v_1^-, v_2^-, \dots\} \quad \text{s.t.} \quad \forall v^- \in \bar{u}, (u, v^-) \notin E_{\text{train}} \quad (11)$$

where elements of $\bar{u}$ are sampled uniformly at random. Arguably, it is possible to increase the accuracy of our models if we sub-sample frequent nodes (i.e. with a high degree) as recommended by [19], however we leave this as future work. In our training loop, we uniformly sample a subset of size K from $\bar{u}$, where K is a hyper-parameter for Negative Sampling. We use a fixed $K = 5$ for all our experiments.

4.4 Negative Sampling

We define an objective, $\mathcal{L}$, that can be computed in linear-time using negative sampling, (similar to [19, Section 2.2]). $\mathcal{L}$ approximates

our quadratic graph likelihood (8), defined as:

$$\mathcal{L} = \mathbb{E}_{(u,v) \sim \mathcal{D}/Z} \left[\log \sigma(g(u,v)) + \sum_{v^- \in \text{Sample}(K, \bar{u})} \log(1 - \sigma(g(u, v^-))) \right], \quad (12)$$

where $\text{Sample}(K, \bar{u})$ uniformly samples K negatives from $\bar{u}$ without replacement and Z is a normalizing constant. Note that the outer expectation $(u, v) \in \mathcal{D}$ is linear and the inner summation goes over K items. We use TensorFlow [28] to obtain the gradients $\frac{\partial \mathcal{L}}{\partial Y_u}, \frac{\partial \mathcal{L}}{\partial \theta}, \frac{\partial \mathcal{L}}{\partial L}, \frac{\partial \mathcal{L}}{\partial R}$ for each mini-batch. We use PercentDelta [1] to optimize all parameters θ, f, g, Y . We only update the anchor embeddings Y_u during the gradient steps on the objective $\mathcal{L}$, as preliminary experiments showed that we get better performance.

5 EXPERIMENTS

For all of our experiments, we simulated $n = 80$ walks from every node, each walk is of length $\tau = 100$, and we used a right and left context window sizes, respectively, for directed and undirected graphs as $(w_l = 0, w_r = 2)$ and $(w_l = 2, w_r = 2)$.

5.1 Datasets

We test our algorithms on directed and undirected graphs. We obtain PPI from [11, 27] and the other datasets from Stanford SNAP [16]. We only use the largest weakly connected component (WCC) from the original graph. The statistics and dataset description are as follows:

Directed graphs:

- (1) **social-epinions**: A social network $|V| = 75,877$ and $|E| = 508,836$. Each directed edge represents whether a user trusts the opinion of another.
- (2) **wiki-vote**: A voting network with $|V| = 7,066$ and $|E| = 103,663$. Nodes are Wikipedia editors. Each directed edges represents a vote that another becomes an administrator.

Undirected graphs:

- (1) **ca-HepTh**: A citation network of High Energy Physics Theory from Arxiv, with $|V| = 17,903$ and $|E| = 197,031$. Each undirected edge represents co-authorship between two author nodes.
- (2) **ca-AstroPh**: A citation network of Astrophysics from Arxiv, with $|V| = 17,903$ and $|E| = 197,031$. Each undirected edge represents co-authorship between two author nodes.
- (3) **PPI**: A protein-protein interaction graph, with $|V| = 3,852$ and $|E| = 20,881$. This is a challenging real-world dataset, where each node is a protein and an edge represents that two proteins interact.
- (4) **ego-Facebook**: A small portion of the Facebook social network, with $|V| = 4,039$ and $|E| = 88,234$. The nodes are users and the edges indicate friendship. We note that this graph is an ego-network graph, which contains only the complete social connections of 10 seed users. Rather than running link-prediction experiments on this graph, we analyze its unique structure through visualization in Section 5.3.

5.2 Link Prediction

We follow the setup in [11] for link prediction. First, given a graph $G = (V, E)$, we partition its edges into two equal size disjoint partitions E_{train} and E_{test} , such that, E_{train} is connected. Second, we sample negative edges for training and testing, E_{train}^- and E_{test}^- , where E_{train}^- is sampled from the compliment of E_{train} and E_{test}^- is sampled from the compliment of E . All train/test edge sets are of equal size. Third, we simulate random walks on E_{train} to get $\mathcal{D}$, using Algorithm 1 and Eq. (9). Fourth, only for directed graphs, we extend E_{test}^- to contain all edges (v, u) s.t. $(u, v) \in E$ and $(v, u) \notin E$. Finally, we train each algorithm and we evaluate ROC-AUC metrics on their ranking of $(E_{\text{test}}, E_{\text{test}}^-)$.

5.2.1 Methods. We report results from various methods, including non-embedding baselines, embedding baselines, and our proposed embedding methods.

Adjacency (non-embedding) Baselines:

These methods require E_{train} during inference. Let $N(u)$ denote the list of node u 's direct neighbors, that are observed according to E_{train} . If the graph is directed, then $N(u)$ only stores outgoing edges. Adjacency baselines score an edge (u, v) as a function of $N(u)$ and $N(v)$. We evaluate against the following:

- (1) **Jaccard Coefficient** models the edge score as

$$g(u, v) = \frac{|N(u) \cap N(v)|}{|N(u) \cup N(v)|}$$

- (2) **Common Neighbors** models the edge score as

$$g(u, v) = |N(u) \cap N(v)|$$

- (3) **Adamic Adar** models the edge score as

$$g(u, v) = \sum_{x \in N(u) \cap N(v)} \frac{1}{\log(|N(x)|)}$$

Embedding Baselines:

These methods use E_{train} to learn embedding Y_u for every graph node u . During inference, they use the learned embedding dictionary Y but *not* the original graph edges. We compare against various state-of-the-art embedding methods.

- (1) Laplacian **EigenMaps** [3] finds the lowest eigenvectors of the graph Laplacian matrix. The eigendecomposition is real iff the Laplacian is symmetric. Therefore, we convert directed graphs to undirected ones during training. During inference, we define the edge scoring function as $g(u, v) = -||Y_u - Y_v||$.
- (2) **node2vec** [11] learns embedding by simulating random walks on E_{train} and minimizing the skipgram objective (Equation 3). We use the author's code to learn node embeddings, then we calculate the hadamard product $Y_u \odot Y_v$ for all node pairs (u, v) in E_{train} or E_{train}^- . Finally, we model an edge score as $g(u, v) = w^T(Y_u \odot Y_v)$ where w is trained using off-the-shelf binary classification algorithm, scikit-learn's Logistic Regression. We train w so that the logistic $\sigma(w^T(Y_u \odot Y_v)) \approx 1$ if $(u, v) \in E_{\text{train}}$ and ≈ 0 if $(u, v) \in E_{\text{train}}^-$. According to our understanding, this is similar to how node2vec performed link prediction [11].
- (3) **DNGR** [5] learns a non-linear (i.e. deep) node embeddings by passing "smoothed" adjacency matrix through a deep auto-encoder. The "smoothing" (called Random Surfing in

	Dataset	Adjacency Methods			Embedding Methods								
		Non-Embedding baselines			Embedding Baselines				Ours: (end-to-end) Graph Likelihood				
		Jaccard	Common Neighbors	Adamic Adar	d	Eigen Maps	node2vec	DNGR	Symmetric		Asymmetric		% Error Reduction
								shallow	deep	shallow	deep		
directed	soc-epinions	0.649	0.649	0.647	8	†	0.725	†	0.694	0.665	0.695	0.825	36.5%
					16	†	0.726	†	0.710	0.713	0.699	0.840	41.4%
					32	†	0.714	†	0.740	0.713	0.700	0.845	45.9%
					64	†	0.699	†	0.766	0.722	0.698	0.834	44.9%
					128	†	0.691	†	0.782	0.743	0.718	0.828	44.5%
	wiki-vote	0.579	0.580	0.562	8	0.613	0.643	0.630	0.603	0.602	0.608	0.871	63.7%
					16	0.607	0.642	0.622	0.623	0.639	0.643	0.900	71.9%
					32	0.600	0.641	0.619	0.642	0.661	0.683	0.911	75.2%
					64	0.613	0.642	0.598	0.660	0.672	0.702	0.917	76.7%
					128	0.622	0.643	0.554	0.682	0.685	0.730	0.917	76.8%
undirected	ca-HepTh	0.765	0.765	0.765	8	0.786	0.731	0.706	0.855	0.848	0.605	0.879	43.2%
					16	0.790	0.787	0.780	0.894	0.826	0.885	0.899	51.9%
					32	0.795	0.858	0.829	0.896	0.886	0.884	0.911	37.8%
					64	0.802	0.886	0.868	0.878	0.884	0.870	0.910	21.3%
					128	0.812	0.901	0.897	0.891	0.897	0.820	0.916	14.6%
	ca-AstroPh	0.942	0.942	0.944	8	0.825	0.811	0.852	0.923	0.925	0.592	0.917	44.1%
					16	0.825	0.833	0.877	0.950	0.923	0.657	0.945	55.8%
					32	0.825	0.899	0.917	0.955	0.938	0.942	0.955	46.1%
					64	0.824	0.934	0.939	0.948	0.936	0.936	0.958	30.7%
					128	0.829	0.955	0.968	0.953	0.936	0.939	0.957	n/a
	PPI	0.766	0.776	0.779	8	0.710	0.733	0.583	0.746	0.763	0.550	0.804	26.6%
					16	0.711	0.707	0.687	0.780	0.772	0.786	0.817	36.7%
					32	0.709	0.691	0.741	0.779	0.784	0.794	0.833	35.5%
					64	0.707	0.671	0.767	0.791	0.767	0.813	0.837	30.0%
					128	0.737	0.698	0.769	0.795	0.787	0.799	0.841	31.0%

Table 1: Link Prediction results from ranking E_{test} across five graph datasets. Numbers shown are the ROC-AUC. Row-wise: the first two datasets are directed and the last four are undirected graphs. Column-wise: The first three methods are adjacency (non-embedding) methods that use the direct neighbors $N(u)$ and $N(v)$ for computing $g(u, v)$. We then list all embedding methods, preceded by d , the dimensionality of the learned embeddings. We report three embedding baselines followed by our four embedding methods. We compare embedding methods across different dimensionality $\{8, 16, 32, 64, 128\}$, marking in bold the top performer for the given dimensionality. For asymmetric embedding methods, we train with half of the dimensionality ($= \{4, 8, 16, 32, 64\}$) since in practice we need to store both sides of the embedding to compute an edge score, therefore every row contains the same total dimensions per node. The last column shows the relative reduction in error of our Asymmetric Deep model compared the best baseline. † indicates that the algorithm runs out of memory on our machine with 32 GB ram.

[5]) is their proposed alternative to random walks, which effectively has a different context weighing from node2vec. We use the author’s code to train the auto-encoder on the adjacency matrix that corresponds to E_{train} . To test different embedding sizes, we only change the size of the last bottleneck layer in their code and keep the remainder of default architecture. We then output the bottleneck layer values for all nodes to and use them for the link prediction task, with scoring function $g(u, v) = Y_u^T Y_v$.

Our Methods:

- (1) **Symmetric Shallow:** $w^T(Y_u \odot Y_v)$. From a modelling perspective, this is identical to node2vec’s model. However, we train this model on our objective (Equation 12) rather than the skipgram objective (Equation 3).
- (2) **Symmetric Deep:** $w^T(f(Y_u) \odot f(Y_v))$. Similar to above, except that the embedding representation is deep.

- (3) **Asymmetric Shallow:** $Y_u^T \times L \times R \times Y_v$. Applying asymmetry directly on the embeddings without a DNN.
- (4) **Asymmetric Deep:** $f(Y_u)^T \times L \times R \times f(Y_v)$. Our full asymmetric formulation, when composed of a single affine projection. For both of our asymmetric methods, after training Y, f_θ, g , we use only the b -dimensional edge representations for inference ($L^T Y_u$ and $R Y_v$).

We train all our models on $\mathcal{D}$ using PercentDelta [1], with learning rate 0.001 and L2 regularization of 0.0001 on all trainable parameters. We find that PercentDelta is crucial, especially that gradients w.r.t. embeddings were large with our initialization scheme. We do model selection using E_{train} and E_{train}^- .

5.2.2 Link Prediction Results. Here we discuss the results of our link prediction experiments, which are presented in Table 1.

Figure 3: Visualization of 2-dimensional embeddings learned from our Asymmetric Deep model on the Facebook dataset. Figures (a-d) render nodes in their original learned coordinates (mapped to pixel space as $(L^T \times f(Y_u)) * \text{radius} + \text{center}$). Figs (a) and (b), respectively render the left- and right-embeddings of all nodes. Each node is plotted as a circle, with an overlay tick whose length is proportional to its degree. Fig (c) combines the left- and right-embedding spaces, dropping the degree ticks for clarity. Fig (d) shows a few selected nodes from the right-embedding, and for each selected node, we draw all its edges to nodes onto the left-embedding.

Figure 4: Unconstrained Asymmetric Embedding Visualization for PPI. Figs (a) and (b) show the left- and right-embedding spaces. Fig (c) shows the combined plot. Fig (d) shows two selected nodes from the right space and all of their edges onto the left space. The 90-degree cone per selected node, highlighting the area where the model has encoded probable edges (i.e. where the node has positive dot-product with the left embedding space). The color of each node in the right embedding space depend on its (x, y) position. The colors of the left embedding space is set to the normalized adjacency matrix multiplied by the colors of the right embedding space. Axis lines $x = 0$ and $y = 0$ are shown in black, $y = x$ and $y = -x$ are shown in grey.

First we turn our attention to directed graphs, where there is a dramatic increase in performance. Specifically on graph socpinions, the asymmetric deep model reduces error over the baseline by 44.9% for 64-dimensional representations. We see that the asymmetric deep representations make more efficient use of their allocated space since its performance at lower dimensions (e.g. $d = 8$) is approximately equal to its performance at higher dimensions (e.g. 64 or 128). The second directed graph, wiki-vote, shows an even stronger performance increase, with a reduction in error of up to 76.8% over the state-of-the-art baselines.

Next, we consider undirected graphs. On the citation networks, ca-HepTh and ca-AstroPh, we see that deep asymmetric approaches

offer large improvements over the baseline when $d = 16$ (respectively, 51.9% and 55.8%), but this lead narrows as the baseline representations are allowed more capacity. Finally, we examine results on the protein-protein interaction network. This is perhaps our most challenging undirected graph, derived from a problem of significant scientific interest. On this network, we observe a large boost from the deep asymmetric model over the baseline method, of up to a 36.7% relative reduction in error ($d = 64$). We note that even when using representations which are 16 times smaller ($d = 8$), the deep asymmetric model has AUC of 0.804, which is a 15% error reduction over the best DNGR baseline ($d = 128$, with AUC of 0.769).

We comment briefly on other observations from methods which optimize the graph likelihood. First, from a modeling prospective,

our shallow symmetric formulation is identical to node2vec's, but they differ in the training objective. This verifies that our proposed graph likelihood produces embeddings that better preserve the graph structure than the Skip-gram objective (Equation 3). Second, shallow models tend to perform much worse in the presence of limited representation size. This is unsurprising, as the a shallow model has to represent each node individually, rather than learning a common latent feature space which can implicitly learn correlations in the data. Third, using asymmetry alone (without a deep model) does not offer nearly as much performance improvement as the asymmetric deep model.

5.3 Manifold Visualizations

We visualize asymmetric embeddings learned for link prediction for two graphs: ego-Facebook network and the PPI network [27]. We train both to be 2 dimensional ($b = 2$), so that we can plot them on this paper without using an external embedding visualization algorithms such as t-SNE. To give two flavors of visualizations, we constraint the embedding of the former to be *circular* but we put no constraints on the latter.

5.3.1 Circular Visualization of ego-Facebook. Similar to our link-prediction setup, we train on half of the edges (i.e. on E_{train}). However, we show on the visualization edges from both partitions E_{train} and E_{test} . In order to easily display the node degrees on the visualization, we constraint the embeddings to be circular (unit-norm) on both sides, specifically as: $\|L^T f(Y_u)\|_2 = \|Rf(Y_v)\|_2 = 1$. Figure 3 shows embeddings. For and setting L2-norm constraints on both sides of the asymmetric embeddings, specifically: $\|L^T f(Y_u)\|_2 = \|Rf(Y_v)\|_2 = 1$. Under this constraint, those 2-dimensional embeddings have only one degree of freedom e.g. angle, and it is straight-forward to show that

$$\arg \max_v \langle L^T f(Y_u), Rf(Y_v) \rangle = \arg \min_v \|L^T f(Y_u) - Rf(Y_v)\|.$$

The visualization shows the two asymmetric embedding spaces are almost disjoint, where nodes from the right space are placed closer to their neighbors in the left space. We also note that the high-degree nodes are closest to the other embedding space. In fact, the highest degree nodes “pull” the left embedding space, as they live within it.

5.3.2 Unconstrained Visualization of PPI. We show in Figure 4 the left- and right- embedding spaces learned for PPI when it is 2-dimensional ($b = 2$). The right-embedding space was colored deterministically. In the right embedding space, the color of node u is based on its right-embedding $(Rf(Y_u)) \in \mathbb{R}^2$ coordinates. In the left embedding space, the color of node v is set to the average of right-embedding colors of v 's neighbors. We see that nodes within 90 degrees have similar colors, showing that our method is embedding the nodes in appropriate positions across the two spaces, to preserve the graph structure.

5.4 Improved Generalization

Most machine learning models are prone to overfitting, showing higher performance metrics on the “train” partition than on the “test” partition. Here we consider an empirical evaluation of our proposed model's overfitting. Specifically, we compute the ratio of

test-over-train accuracy. If this test-over-train ratio = 1, it means that the model does as well on (held-out) test data as it does on the training data. We note that this frequently does not occur in practice – typically models overfit the training data, and the ratio is < 1 . Nonetheless, Table 2 shows that adding DNN $f()$ brings this ratio closer to 1. For example, on wiki-vote and soc-epinions, using the DNN $f()$ increases this ratio by over 4%. Since we average this ratio across all our runs, we report the t-test numbers concluding that all our numbers are statistically significant ($p < 0.01$).

In Table 3 we show that adding $f()$ can be seen as a regularization on the embeddings. In other words, deeper models produce $f(Y)$ with *consistent* embedding L2-norms, across all datasets and across all runs, while shallow models produce Y with a wider variation of L2-norms.

5.5 Parameter Sensitivity

In order to understand the impact of the representation size as a function of task performance, we varied the number of dimensions in the model from $d = 4$ to $d = 512$. The results of this experiment are shown in Figure 5.

6 DISCUSSION

We have proposed a “deep” asymmetric model which learns Y jointly with f and g as $Y \rightarrow f \rightarrow g$. One might think that a “shallow” asymmetric model $Y \rightarrow g$ should be able to learn Y , identically to how a deeper counter-part learns $f(Y)$. However this is not necessarily the case. In this section we motivate why using the DNN $f()$ is helpful and reference empirical evidence that supports our claims.

First, $f()$ removes degrees of freedom as it passes an embedding Y_u through the DNN activation functions. Here, $f()$ can be seen as a regularizer over the embeddings. In particular, the output of $f(Y_u)$ is *bounded*, as the last BatchNorm layer [14] ensures that the $f(Y_u)$ has approximately a fixed mean and fixed variance across batches. On the other hand, shallow models can learn an *unbounded* Y_u . We summarize the embedding norm statistics on real datasets in Table 3.

Second, as $f()$ can constrain the embeddings with less degrees of freedom, this reduces overfitting and improves generalization. Table 2 shows that deeper models have higher test-over-train accuracy metrics. This generalization is important since our proposed objective (the graph likelihood) is not directly a link prediction objective and the evaluation data ($E_{\text{test}}, E_{\text{test}}^-$) is not observed during training.

Finally, the hidden layers in $f()$ find correlations in the data, as many graph nodes have similar connections. It finds a smaller non-linear dimensional space that the nodes live in. For some D -dimensional embedding $Y \in \mathbb{R}^{|V| \times D}$ neural network can map Y onto a lower dimensional manifold as $f : \mathbb{R}^D \rightarrow \mathbb{R}^d$ where $d < D$. In fact, our experiments in Table 1 show that our deep model performs quite well when provided with very few dimensions.

7 RELATED WORK

There is a rapidly growing body of literature on applying neural networks to problems which have as input a graph. We divide the related work into two broad groups, based on whether it concerns

Figure 5: ROC-AUC results for our models versus the baselines measured on a suite of datasets.

Dataset	mean($\frac{\text{test AUC}}{\text{train AUC}}$)		Statistical Significance	
	shallow asymmetric	deep asymmetric	t-statistic	p-value
soc-epinions	0.841	0.882	5.797673	1.53E-06
wiki-vote	0.908	0.948	3.881161	4.32E-04
ca-HepTh	0.881	0.915	5.202880	1.62E-05
ca-AstroPh	0.946	0.970	5.946066	4.08E-07
ppi	0.865	0.893	4.187474	8.45E-05

Table 2: Showing generalization performance by averaging ($\frac{\text{test AUC}}{\text{train AUC}}$) across all runs for all datasets under two settings: “shallow” VS “deep” asymmetric i.e. absence VS presence of DNN $f(\cdot)$. Last two columns show the t-test for the difference of test-over-train AUC between the two settings. Adding a DNN to the model is a statistically significant improvement on all graphs with $p < 0.001$.

Dataset	Shallow Symmetric			Shallow Asymmetric			Deep Symmetric			Deep Asymmetric		
	25 th	50 th	75 th	25 th	50 th	75 th	25 th	50 th	75 th	25 th	50 th	75 th
wiki-vote	0.106	0.202	0.327	0.122	0.142	0.152	0.597	0.901	1.119	0.811	1.096	1.938
soc-epinions	0.382	0.526	0.754	0.299	0.345	0.430	0.888	1.147	1.404	1.276	1.881	3.590
ppi	1.884	4.593	7.842	0.858	1.197	3.801	0.825	1.095	1.410	1.015	1.443	2.645
ca-HepTh	7.370	10.871	12.426	3.093	3.916	7.892	0.957	1.364	1.709	1.120	1.417	2.292
ca-AstroPh	12.069	326.483	4282.095	1.108	4.648	24.271	0.874	1.273	1.999	1.065	1.826	3.062

Table 3: Inner-quartile Ranges of Standard Deviations of Embedding Norm, across all runs. For shallow models, we calculate $\text{std}_{u \in V}(\|Y_u\|)$, where $\text{std}_{u \in V}(\cdot)$ is the standard deviation for all $u \in V$, and we display the statistics across all runs (e.g. different embedding dimensions). For deep models, we calculate $\text{std}_{u \in V}(\|f(Y_u)\|)$

graph classification (e.g. assigning a label to G), or learns a representation per node for preserving the graph structure.

Discriminative Learning on Graph-Structured Data. These algorithms learn representations (at node/edge/graph) that are used for a discriminative classification task (per node/edge/graph). These methods are powerful for discrimination but they strictly rely on the graph structure as “golden ground-truth” to propagate information – e.g. Graph-convolutional methods, using adjacency edges to define non-Euclidean patches [2, 4, 20] and some operate in the fourier domain [4, 8, 13]. In addition, some discriminative representations include fixed-point methods, recursively defining node features as a function of its neighbors by “unrolling a few steps” [9] or until fixed-point convergence is reached [7, 10, 17, 25]. We differ from all these methods, since they receive an external loss (e.g. label) and assume that the graph is completely observed. For example conditional independence assumptions made by the Markov Models of [7] explicitly use the graph structure. Unlike our work, these

methods have no obvious way to estimate the score/probability of an edge, as the existence of the edge was inherently used to pass discriminative information.

Structure-Preserving Embeddings. These methods learn one embedding per graph node, with an objective that maximizes (/ minimizes) the product (/ distance) of node embeddings if they are neighbors in the input graph. They are most related to our work. In fact, our work builds on the approach introduced by Deepwalk [22], which learns node embeddings using simulated random walks. These node embeddings have been used as features for various tasks on networks, such as node classification [22], user profiling [24], and link prediction [11]. Some extensions of Deepwalk include: Walklets [23], skipping nodes in the random walk to discover hierarchical structure; node2vec [11], parameterizing the random walk process to allow more focused discovery of structural relationships; author2vec [15], augmenting nodes with bag-of-words; and

Tri-Party DNN [21], modeling heterogeneous graphs with three different node types. Other node-centric methods are concerned with shorter dependencies in the graph [29, 30]. Finally, meta-embedding approaches, such as HARP [6], have been proposed as general methods for improving node representations.

Our work differs from existing random walk methods in three ways. First, we explicitly model asymmetric relationships between nodes. Even though random walk methods we surveyed **respect** edge direction during the walk, they do **not** model edge direction and represent (u, v) identically to (v, u) [11, 15, 18, 22, 23]. This flexibility better models the heterogeneity which occurs in real world networks, where social relationships may not be reciprocal (i.e. the graph is *directed*), or typically have a very unbalanced degree distribution. Second, our node embeddings are produced by a deep neural network (DNN), unlike the shallow (effectively 1-layer) networks previously used. Third, rather than a 2-stage optimization of first training embeddings on random walk sequences, followed by learning a task-specific classifier, we propose a graph likelihood and use it to jointly train the embeddings, the manifold-mapping DNN, and the edge function. Even though the graph likelihood does *not* match a link-prediction loss, it produces superior results on link-prediction using fewer dimensions.

8 CONCLUSION

We introduced a novel method for integrating directed edge information for learning continuous representation for graphs. Our method *explicitly models edges* as functions of node representations. We optimize this model using a new objective function, the *graph likelihood*, which we use to jointly learn the edge function and node representations.

Our empirical evaluation focused on link prediction tasks using a number of graphs collected from real world applications. Our experimental results show that our proposed objective is better than the skipgram objective even when the model are identical. Our results also show that modeling edges as asymmetric affine projections through the node representation space, helps produce more accurate and compact embedding spaces. In particular, we show that asymmetric edge modeling, when trained with our objective, improves performance over state-of-the-art approaches, especially on directed graphs, reducing error by up to $\approx 70\%$ and $\approx 50\%$, respectively, on directed and undirected graphs while using the same number of dimensions per node. We create visualizations to interpret the learned representations, showing that our method embeds nodes in appropriate places along the embedding manifold.

In addition to AUC metric improvements, explicit edge modeling allows us to learn smaller embeddings. The representations learned through our model are more efficient at utilizing the available space. Our embeddings are able to outperform the baseline even when outputting 8x fewer dimensions per node. We believe that explicitly modeling edge representations addresses a substantial problem in the related work, and can enable many avenues of future investigation for learning continuous representation of graphs.

REFERENCES

- [1] Sami Abu-El-Hajja. 2017. Proportionate gradient updates with PercentDelta. In *arXiv*.
- [2] James Atwood and Don Towsley. 2016. Diffusion-Convolutional Neural Networks. In *Advances in Neural Information Processing Systems (NIPS)*.
- [3] M. Belkin and P. Niyogi. 2001. Laplacian eigenmaps and spectral techniques for embedding and clustering. In *Advances in Neural Information Processing Systems (NIPS)*.
- [4] Joan Bruna, Wojciech Zaremba, Arthur Szlam, and Yann LeCun. 2013. Spectral networks and deep locally connected networks on graphs. In *International Conference on Learning Representations*.
- [5] Shaosheng Cao, Wei Lu, and Qionghai Xu. 2016. Deep Neural Networks for Learning Graph Representations. In *Proceedings of the Association for the Advancement of Artificial Intelligence*.
- [6] Haochen Chen, Bryan Perozzi, Yifan Hu, and Steven Skiena. 2017. HARP: Hierarchical Representation Learning for Networks. *arXiv preprint arXiv:1706.07845* (2017).
- [7] Hanjun Dai, Bo Dai, and Le Song. 2016. Discriminative Embeddings of Latent Variable Models for Structured Data. In *International Conference on Machine Learning (ICML)*.
- [8] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. 2016. Convolutional Neural Networks on Graphs with Fast Localized Spectral Filtering. In *Advances in Neural Information Processing Systems (NIPS)*.
- [9] D. Duvenaud, D. Maclaurin, J. Iparraguirre, R. Bombarell, T. Hirzel, A. Aspuru-Guzik, and R. Adams. 2015. Convolutional Networks on Graphs for Learning Molecular Fingerprints. In *Advances in Neural Information Processing Systems (NIPS)*.
- [10] M. Gori, G. Monfardini, and F. Scarselli. 2005. A new model for learning in graph domains. In *Proc. International Joint Conference on Neural Networks (IJCNN)*.
- [11] A. Grover and J. Leskovec. 2016. node2vec: Scalable Feature Learning for Networks. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- [12] L. Hagen and A. Kahng. 1992. New spectral methods for ratio cut partitioning and clustering. In *IEEE Trans. Computer-Aided Design*.
- [13] Mikael Henaff, Joan Bruna, and Yann LeCun. 2015. Deep Convolutional Networks on Graph-Structured Data. In *arXiv:1506.05163*.
- [14] Sergey Ioffe and Christian Szegedy. 2015. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. *Journal of Machine Learning Research (JMLR)*.
- [15] G. J. S. Ganguly, M. Gupta, V. Varma, and V. Pudi. 2016. Author2Vec: Learning Author Representations by Combining Content and Link Information. In *Proceedings of the 25th International Conference Companion on World Wide Web (WWW '16 Companion)*.
- [16] J. Leskovec and A. Krevl. 2014. SNAP Datasets: Stanford large network dataset collection. <http://snap.stanford.edu/data>
- [17] Y. Li, D. Tarlow, M. Brockschmidt, and R. Zemel. 2016. Gated Graph Sequence Neural Networks. In *International Conference on Learning Representations*.
- [18] Y. Luo, Q. Wang, B. Wang, and L. Guo. 2015. Context-Dependent Knowledge Graph Embedding. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- [19] T. Mikolov, I. Sutskever, K. Chen, G. Corrado, and J. Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems*.
- [20] M. Niepert, M. Ahmed, and K. Kutzkov. 2016. Learning Convolutional Neural Networks for Graphs. In *International Conference on Machine Learning (ICML)*.
- [21] S. Pan, J. Wu, X. Zhu, C. Zhang, and Y. Wang. 2016. Tri-Party Deep Network Representation. In *International Joint Conference on Artificial Intelligence*.
- [22] B. Perozzi, R. Al-Rfou, and S. Skiena. 2014. DeepWalk: Online Learning of Social Representations. In *Knowledge Discovery and Data Mining*.
- [23] B. Perozzi, V. Kulkarni, H. Chen, and S. Skiena. 2017. Don't Walk, Skip! Online Learning of Multi-scale Network Embeddings. In *2017 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*.
- [24] Bryan Perozzi and Steven Skiena. 2015. Exact Age Prediction in Social Networks. (2015), 2 pages.
- [25] F. Scarselli, M. Gori, A.C. Tsoi, M. Hagenbuchner, and G. Monfardini. 2009. The Graph Neural Network Model. In *IEEE Trans. on Neural Networks*.
- [26] J. Shi and J. Malik. 2000. Normalized cuts and image segmentation. In *IEEE Trans. Pattern Anal. Mach. Intell.*
- [27] C. Stark, B.J. Breitkreutz, T. Reguly, L. Boucher, A. Breitkreutz, and M. Tyers. 2006. BioGRID: A General Repository for Interaction Datasets. In *Nucleic Acids Research*. <https://www.ncbi.nlm.nih.gov/pubmed/16381927>
- [28] TensorflowTeam. 2015. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. <http://tensorflow.org/> Software available from tensorflow.org.
- [29] D. Wang, P. Cui, and W. Zhu. 2016. Structural Deep Network Embedding. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- [30] H. Wang, X. Shi, and D-Y. Yeung. 2017. Relational deep learning: A deep latent variable model for link prediction. In *Conference on Artificial Intelligence (AAAI)*.
- [31] Y. Yu, C. Fang, and Z. Liao. 2015. Piecewise Flat Embedding for Image Segmentation. In *IEEE International Conference on Computer Vision (ICCV)*.