# Constructing Graph Node Embeddings via Discrimination of Similarity Distributions

Stanislav Tsepa
*Skolkovo Institute of Science and Technology*
(Skoltech)
*Moscow Institute of Physics and Technology*
*Institute for Information Transmission Problems of RAS*
Moscow, Russia
s.tsepa@skoltech.ru

Maxim Panov
*Skolkovo Institute of Science and Technology*
(Skoltech)
*Institute for Information Transmission Problems of RAS*
Moscow, Russia
m.panov@skoltech.ru

arXiv:1810.03032v1 [stat.ML] 6 Oct 2018

***Abstract*—The problem of unsupervised learning node embeddings in graphs is one of the important directions in modern network science. In this work we propose a novel framework, which is aimed to find embeddings by *discriminating distributions of similarities (DDoS)* between nodes in the graph. The general idea is implemented by maximizing the *earth mover distance* between distributions of decoded similarities of similar and dissimilar nodes. The resulting algorithm generates embeddings which give a state-of-the-art performance in the problem of link prediction in real-world graphs.**

***Index Terms*—Graph node embeddings, representation learning, Wasserstein distance, unsupervised learning, link prediction**

## I. Introduction

The quality of machine learning methods largely depends on the particular representation (or features) chosen for the data. The majority of modern machine learning methods work with the objects represented as numerical vectors. In some problems, such as object recognition on images and video, speech recognition and natural language processing, the initial feature space formally can be considered as a vector space. The obstacle is that it has a complex structure and a very high dimension which requires the construction of methods to transform the original representation into a more concise and informative one. Problems with objects of a discrete nature (in particular, with graphs) also require that informative continuous numerical representations have to be found. Recently network representation learning attracted a lot of attention which lead the development of many new methods (e.g., see recent reviews [3], [8]).

The essence of the network representation learning problem (or embedding problem) is to represent a graph, a subgraph or a node as a point in low-dimensional Euclidean space; in such a form they can be further used in traditional machine learning pipelines. In what follows we will focus on node embeddings. A usual assumption is that nodes can be represented in a space with dimension $d \ll n$, where $n$ is the number of nodes in the graph. The obtained embeddings can be further used for node classification [7], community detection [6], link prediction [1] or visualization [5].

The majority of now existing graph embedding algorithms focus on proximity preserving embeddings, in which the nodes positioned in a close network proximity are considered to be similar and their embedding points should be placed close to each other in embedding space. We base our research on this assumption and, meanwhile, use the area of representation learning for images as source of inspiration. Recently the histogram loss approach was proposed [18], where the embeddings are learned by minimizing a certain distance between inter- and intra-class similarity distributions. We extend their approach to graphs by introducing a more suitable distance between the distributions which is inspired by a Wasserstein distance. We demonstrate the efficiency of the proposed approach on a series of experiments with real-world graphs.

## II. Related work

Currently the methods of learning mode representations in graphs are rapidly developing. The earlier approaches are based on classical dimension reduction and on finding node embeddings via matrix factorization [2], [14]. Matrix factorization can be really time consuming for large graphs, and a usual way to speed-up embedding learning is to use random walks over the graph. The idea is to find such vector representations of vertices which can describe well the probabilities of particular vertex sequences in these walks. This idea is the base for the whole family of methods, including DeepWalk [15], LINE [16] and node2vec [7]. Note that all these methods use internally the algorithm word2vec [12], which is based on the optimization of the logistic-like loss and was initially proposed for word embeddings. Further developments of random walk-based methods mainly focus on various schemes of random walk through the graph, which allow them to take into account various structural features of graphs. Recently, many attempts were made to adopt neural networks to graph-structured data (see [4], [19] among many others). In the next section we specify the considered framework and further discuss some related embedding learning approaches.

The research was supported by the Russian Science Foundation grant (project 14-50-00150).

## III. Learning graph embeddings in unsupervised way

We start by formalizing the problem of node representation learning in graphs following [8]. We observe an undirected and unweighted graph $\mathcal{G} = (\mathcal{V}, E)$ with $n = |V|$ vertices and denote by $\mathbf{A} \in \{0,1\}^{n\times n}$ the adjacency matrix of graph $\mathcal{G}$ with $n$ nodes. We also assume that certain similarity matrix is given $\mathbf{S} = \left[s_{ij}\right]_{i,j=1}^{n}$, where the value $s_{ij}$ determines how similar the nodes $i$ and $j$ are, and, respectively, how close their embeddings should be. The considered framework consists of several important parts:

1) encoder function which maps nodes to the latent representations
$$\mathcal{E}\colon \mathcal{V} \to \mathbb{R}^d;$$
2) decoder function which maps pairs of node embeddings to node proximity measure
$$\mathcal{D}\colon \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R};$$
3) loss function $\mathcal{L}$ which measures how close reconstructed proximity values $\mathcal{D}(\mathcal{E}_i, \mathcal{E}_j)$ are to the corresponding reference values $s_{ij}$.

In what follows we discuss different types of encoders and loss functions as the choice of them largely distinguishes modern embedding learning algorithms.

### A. Loss functions

We describe two important types of losses considered in the literature.

*1) Pairwise loss:* in this case, the goal of optimization is to minimize the sum of reconstruction errors for pairwise similarities of nodes $\ell\big(\mathcal{D}(\mathcal{E}_i, \mathcal{E}_j), s_{ij}\big)$:

$$\mathcal{L}(\Theta) = \sum_{(i,j)\in\Omega} \ell\big(\mathcal{D}(\mathcal{E}_i, \mathcal{E}_j), s_{ij}\big),$$

where decoder is usually considered to be non-parametrized and optimization is done over the parameters of encoder $\Theta$.

*2) Nodewise autoencoder loss:* this approach was recently proposed by [17], where the authors consider the similarity of node embeddings as distributions:

$$\mathcal{L}(\Theta) = \sum_{i=1}^{n} \ell\big(\mathcal{D}(\mathcal{E}_i, \cdot), \mathbf{s}_i\big),$$

where $\mathcal{D}(\mathcal{E}_i, \cdot)$ and $\mathbf{s}_i$ are vectors of decoded similarities and observed similarities of $i$-th node respectively. One of the possible choices for the loss function is Kullback-Leibler divergence. A similar approach is considered in [4], [20], where autoencoders were constructed considering the classical $L_2$ loss.

### B. Encoders

The standard approach is the so-called embedding lookup, where

$$\mathcal{E}_i = \mathbf{Z}\mathbf{e}_i, \tag{1}$$

where $\mathbf{Z} \in \mathbb{R}^{d\times n}$ is an embedding matrix and $\mathbf{e}_i$ is a vector with 1 on the $i$-th place and zeros elsewhere.

The other possible approach is to treat the local neighborhood of the node as a feature vector and consider

$$\mathcal{E}_i = f(\mathbf{a}_i), \tag{2}$$

where $\mathbf{a}_i$ is the $i$-th column of the adjacency matrix $\mathbf{A}$ and $f\colon \mathbb{R}^n \to \mathbb{R}^d$ is some function. The particular choice of function $f$ might be a neural network [4], [20].

Also one can consider linear function $f$:

$$\mathcal{E}_i = \mathbf{W}\mathbf{a}_i + \mathbf{b}, \tag{3}$$

where $\mathbf{W} \in \mathbb{R}^{d\times n}$ is a parameter matrix, $\mathbf{b} \in \mathbb{R}^d$ is an intercept vector and a full set of encoder parameters is $\Theta = (\mathbf{W}, \mathbf{b})$. If adjacency matrix $\mathbf{A}$ has rank at least $d$ then the expressive ability of such an encoder is exactly equal to the one in a direct embedding approach (1).

## IV. Discrimination of Similarity Distributions

In this paper we introduce a different approach which considers a whole set of similarities between nodes and works out a discriminative loss between the distributions of $\mathcal{D}(\mathcal{E}_i, \mathcal{E}_j)$ in the pairs of similar and non similar nodes. Currently, this approach assumes that the graph is sparse so that similarity matrix $\mathbf{S} = \left[s_{ij}\right]_{i,j=1}^{n}$ is also sparse (we consider either adjacency matrix $\mathbf{S} = \mathbf{A}$ or second order proximities $\mathbf{S} = \mathbf{A}^2$). Consider the set of all positive similarities $L^+ = \{\mathcal{D}(\mathcal{E}_i, \mathcal{E}_j)\colon s_{ij} > 0\}$ and the set of all pairs of nodes with zero similarity $L^- = \{\mathcal{D}(\mathcal{E}_i, \mathcal{E}_j)\colon s_{ij} = 0\}$.

Our main assumption is that the embedding should allow us to distinguish between similar and non similar nodes. In particular, decoded similarities should be higher for similar nodes. If we treat respective positive similarity $s_{ij}$ as a weight for the considered decoded similarity value $\mathcal{D}(\mathcal{E}_i, \mathcal{E}_j)$, then we consider distributions $P^+ = P^+(\Theta)$ and $P^- = P^-(\Theta)$ of decoded similarities in $L^+$ and $L^-$, respectively, and define the loss function as

$$\mathcal{L}(\Theta) = -D\big(P^+(\Theta), P^-(\Theta)\big), \tag{4}$$

where $D$ is a distance between distributions. It might be KL-divergence, Hellinger distance or, for example, Wasserstein distance. Thus, our aim is to maximize the distance between distributions of positive and negative pairs.

We are going to proceed with linear encoder (3) and suggest to use Pearson correlation as a decoder function $\mathcal{D}$:

$$\mathcal{D}(\mathcal{E}_i, \mathcal{E}_j) = \frac{\mathcal{E}_i^{\mathrm{T}}\mathcal{E}_j}{\|\mathcal{E}_i\|\|\mathcal{E}_j\|}.$$

In this case $\mathcal{D}(\mathcal{E}_i, \mathcal{E}_j) \in [-1, 1]$ which is convenient for our purposes.

For the implementation of discriminative loss (4) we follow [18] and approximate distributions of decoded similarities in $L^+$ and $L^-$ by histogram estimators $P^+$ and $P^-$ with linear slope in each bin.

As a distance between distributions we suggest to use 1-D Wasserstein distance (also known as an *"earth mover distance"*, *EMD*) which for the histogram case can be computed as [11]:

$$EMD(P^+, P^-) = \sum_{i=1}^{N_b} |\varphi_i|,$$

where $N_b$ is the number of bins in the histogram and

$$\varphi_i = \sum_{j=1}^{i} \left( \frac{P_j^+}{\|P^+\|_1} - \frac{P_j^-}{\|P^-\|_1} \right).$$

However, in our case it is required that the distribution of similar pairs $P^+$ should be to the right from the distribution of dissimilar pairs $P^-$. The following simple modification allows to avoid unnecessary local optima:

$$EMD_{asym}(P^+, P^-) = \sum_{i=1}^{N_b} \varphi_i. \tag{5}$$

Maximization of modified Wasserstein distance (5) makes distribution $P^+$ concentrating near 1, while distribution $P^-$ concentrates near $-1$. However, it is natural to assume that nodes, which are far away in the graph, should have the embeddings that are independent from each other rather than opposite. That is why we propose to keep in the computation of Wasserstein distance only the part $P^-_{cut}$ of negative distribution $P^-$ corresponding to nonnegative decoded similarities $\mathcal{D}(\mathcal{E}_i, \mathcal{E}_j) >= 0$.

Finally, our objective function becomes

$$\mathcal{L}(\Theta) = EMD_{asym}(P^-_{cut}(\Theta), P^+(\Theta)), \tag{6}$$

which can be optimized via gradient-based methods to find optimal embedding configuration. The gradient of Wasserstein distance between histograms is easily computable, see [11], [18].

## V. Experiments

In this section we will discuss the experimental evaluation of the proposed algorithm. The whole algorithm was implemented in Tensorflow[1], while optimization of the functional (6) over parameters (embedding matrix $\mathbf{W}$ and vector $\mathbf{b}$) was performed via stochastic gradient descent.

The further speedup of the algorithm was achieved by subsampling set $L^-$ (the so-called negative sampling [12]) which is necessary as real-world graphs are usually sparse with $|L^-| \gg |L^+|$.

[1]The code of the algorithm and all the experiments is available at "https://github.com/premolab/GraphEmbeddings.

| Name | Number nodes | Number edges |
|---|---|---|
| Books about US Politics [9] | 105 | 441 |
| American College Football [13] | 115 | 613 |
| Email EU [21] | 986 | 25552 |
| Facebook [10] | 4039 | 88234 |

TABLE I
LIST OF REAL-WORLD NETWORKS USE IN OUR EXPERIMENTS.

### A. Experimental setup

For the experiments we used several real-world networks with number of nodes varying from 100 to 4000. The information about the datasets is summarized in the Table I.

In our experiments we focus on the *link prediction* problem which offers an universal way to estimate the quality of the embedding for any network as it does not require any additional data except for the graph itself. The solution pipeline starts with constructing the embedding based on the part of graph edges and then checks how well the missing edges can be predicted basing on the embeddings. More precisely, the pipeline is as follows:

1) The set of all edges $E$ is randomly divided into two parts: $E_{train}$ and $E_{test}$[2].
2) The embedding $\{\mathcal{E}_v,\ v \in \mathcal{V}\}$ is constructed basing on the subgraph $\mathcal{G}_{train} = (\mathcal{V}, E_{train})$.
3) Define variables

$$y_{uv} = \begin{cases} 1, \text{if } (u,v) \in E, \\ 0, \text{if } (u,v) \notin E, \end{cases}$$
$$X_{uv} = \langle \mathcal{E}_u, \mathcal{E}_v \rangle,$$

   where $\langle \cdot, \cdot \rangle$ means concatenation.
4) Construct classifier (logistic regression in our experiments) basing on the training data set $\{(X_{uv}, y_{uv}), (u,v) \in E_{train}\}$ and estimate the classification quality on the test set $\{(X_{uv}, y_{uv}), (u,v) \in E_{test}\}$.

### B. Results

We compare the results of the proposed algorithm (DDoS) with several state-of-the-art algorithms belonging to different subcategories of embedding algorithms: random walk based algorithms (DeepWalk [15]), direct matrix factorization approaches (HOPE [14]) and a neural network based autoencoders (SDNE [19]).

The results are summarized in Table II. As we can see, in the majority of cases DDoS embeddings allow us to achieve better results than their competitors. Interestingly the usage of the parametrized encoder (3) instead of the embedding lookup (1) has resulted in a much faster convergence of the algorithm, see Figure 1.

[2]In our experiments we set $|E_{train}| = |E_{test}| = \frac{|E|}{2}$.

| Dataset | $d$ | HOPE | DeepWalk | SDNE | DDoS |
|---|---|---|---|---|---|
| Books about US Politics | 4 | 0.89 | 0.87 | 0.78 | **0.90** |
| | 8 | **0.90** | 0.89 | 0.79 | **0.90** |
| | 16 | **0.92** | 0.90 | 0.81 | 0.90 |
| | 32 | **0.93** | 0.90 | 0.75 | 0.90 |
| American College Football | 4 | 0.75 | 0.85 | 0.77 | **0.87** |
| | 8 | 0.86 | **0.90** | 0.82 | **0.90** |
| | 16 | 0.90 | **0.91** | 0.84 | **0.91** |
| | 32 | 0.92 | 0.92 | 0.84 | **0.93** |
| Email EU | 4 | 0.74 | 0.81 | 0.89 | **0.90** |
| | 8 | 0.83 | 0.86 | 0.90 | **0.92** |
| | 16 | 0.90 | 0.89 | 0.92 | **0.93** |
| | 32 | **0.93** | 0.90 | **0.93** | **0.93** |
| Facebook | 4 | 0.84 | 0.94 | 0.96 | **0.97** |
| | 8 | 0.92 | **0.98** | 0.96 | **0.98** |
| | 16 | 0.96 | **0.99** | 0.97 | **0.99** |
| | 32 | 0.97 | 0.99 | 0.98 | **1.00** |

TABLE II
LINK PREDICTION RESULTS (ROC-AUC VALUES).

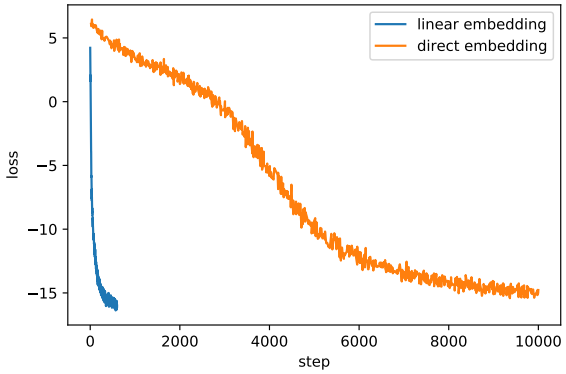


Fig. 1. Steps of loss optimization for the parametrized encoder (3) and the embedding lookup (1) in orange and blue respectively.

## VI. CONCLUSIONS AND OUTLOOK

In this work we propose a simple but powerful approach for constructing the graph embeddings based on discrimination of similarity distributions. We show the way to implement the general idea by using maximization of a specially tuned Wasserstein distance. The series of experiments with the link prediction in real-world graphs convincingly demonstrate the superiority of the proposed approach over its competitors.

Our works offers a number of directions for further development. Among them is the scalability of the algorithm that should be improved first of all. It can be achieved by combining the proposed criterion with sampling graph nodes via random walks. The other direction is to test DDoS embeddings in other network analysis tasks such as community detection and semi-supervised node classification. Extensions to directed and weighted graphs also seem to be of a great interest. Finally, the usage of a non-linear encoder (2) parametrized by a neural network can be a promising direction for further investigation and improvement.

## REFERENCES


[1] Lars Backstrom and Jure Leskovec. Supervised random walks: predicting and recommending links in social networks. In *Proceedings of the fourth ACM international conference on Web search and data mining*, pages 635–644. ACM, 2011.

[2] Mikhail Belkin and Partha Niyogi. Laplacian eigenmaps and spectral techniques for embedding and clustering. In *Advances in neural information processing systems*, pages 585–591, 2002.

[3] Hongyun Cai, Vincent W Zheng, and Kevin Chang. A comprehensive survey of graph embedding: problems, techniques and applications. *IEEE Transactions on Knowledge and Data Engineering*, 2018.

[4] Shaosheng Cao, Wei Lu, and Qiongkai Xu. Deep neural networks for learning graph representations. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, pages 1145–1152. AAAI Press, 2016.

[5] MC Ferreira De Oliveira and Haim Levkowitz. From visual data exploration to visual data mining: a survey. *IEEE Transactions on Visualization and Computer Graphics*, 9(3):378–394, 2003.

[6] Santo Fortunato. Community detection in graphs. *Physics reports*, 486(3-5):75–174, 2010.

[7] Aditya Grover and Jure Leskovec. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 855–864. ACM, 2016.

[8] William L Hamilton, Rex Ying, and Jure Leskovec. Representation learning on graphs: Methods and applications. *arXiv preprint arXiv:1709.05584*, 2017.

[9] Valdis Krebs. Books about us politics, 2004.

[10] Jure Leskovec and Julian J Mcauley. Learning to discover social circles in ego networks. In *Advances in neural information processing systems*, pages 539–547, 2012.

[11] Manuel Martinez, Makarand Tapaswi, and Rainer Stiefelhagen. A Closed-form Gradient for the 1D Earth Mover's Distance for Spectral Deep Learning on Biological Data. In *ICML 2016 Workshop on Computational Biology (CompBio@ICML16)*, Jun. 2016.

[12] 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.

[13] Mark EJ Newman and Michelle Girvan. Finding and evaluating community structure in networks. *Physical review E*, 69(2):026113, 2004.

[14] Mingdong Ou, Peng Cui, Jian Pei, Ziwei Zhang, and Wenwu Zhu. Asymmetric transitivity preserving graph embedding. In *Proceedings of the 22Nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16, pages 1105–1114, New York, NY, USA, 2016. ACM.

[15] Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. Deepwalk: Online learning of social representations. *CoRR*, abs/1403.6652, 2014.

[16] Jian Tang, Meng Qu, Mingzhe Wang, Ming Zhang, Jun Yan, and Qiaozhu Mei. Line: Large-scale information network embedding. In *Proceedings of the 24th International Conference on World Wide Web*, pages 1067–1077. International World Wide Web Conferences Steering Committee, 2015.

[17] Anton Tsitsulin, Davide Mottin, Panagiotis Karras, and Emmanuel Müller. VERSE. In *Proceedings of the 2018 World Wide Web Conference on World Wide Web - WWW '18*, pages 539–548, New York, New York, USA, 2018. ACM Press.

[18] Evgeniya Ustinova and Victor Lempitsky. Learning deep embeddings with histogram loss. In *Advances in Neural Information Processing Systems*, pages 4170–4178, 2016.

[19] Daixin Wang, Peng Cui, and Wenwu Zhu. Structural deep network embedding. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1225–1234. ACM, 2016.

[20] Daixin Wang, Peng Cui, and Wenwu Zhu. Structural Deep Network Embedding. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD '16*, pages 1225–1234, New York, New York, USA, 2016. ACM Press.

[21] Hao Yin, Austin R Benson, Jure Leskovec, and David F Gleich. Local higher-order graph clustering. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 555–564. ACM, 2017.