



Document Network Projection in Pretrained Word Embedding Space

Antoine Gourru^(✉) , Adrien Guille , Julien Velcin , and Julien Jacques 

Université de Lyon, Lyon 2, ERIC EA3083, Lyon, France

{antoine.gourru, adrien.guille, julien.velcin, julien.jacques}@univ-lyon2.fr

Abstract. We present Regularized Linear Embedding (RLE), a novel method that projects a collection of linked documents (e.g., citation network) into a pretrained word embedding space. In addition to the textual content, we leverage a matrix of pairwise similarities providing complementary information (e.g., the network proximity of two documents in a citation graph). We first build a simple word vector average for each document, and we use the similarities to alter this average representation. The document representations can help to solve many information retrieval tasks, such as recommendation, classification and clustering. We demonstrate that our approach outperforms or matches existing document network embedding methods on node classification and link prediction tasks. Furthermore, we show that it helps identifying relevant keywords to describe document classes.

Keywords: Document network embedding · Representation Learning

1 Introduction

Information retrieval methods require relevant compact vector space representations of documents. The classical bag of words cannot capture all the useful semantic information. Representation Learning is a way to go beyond and boost the performances we can expect in many information retrieval tasks [6]. It aims at finding low dimensional and dense representations of high dimensional data such as words [12] and documents [2, 10]. In this latent space, proximity reflects semantic closeness. Many recent methods use those representations for information retrieval tasks: capturing user interest [16], query expansion [9], link prediction and document classification [20].

In addition to the textual information, many corpora include links between documents, such as bibliographic networks (e.g., scientific articles linked with citations or co-authorship) and social networks (e.g., tweets with ReTweet relations). This information can be used to improve the accuracy of document representations. Several recent methods [11, 20] study the embedding of networks with textual attributes associated to the nodes. Most of them learn continuous representations for nodes independently of a word-vector representation. That is to say, documents and words do not *lie* in the same space. It is interesting to find

a common space to represent documents and words when considering many tasks in information retrieval (query expansion) and document analysis (description of document clusters). Our approach allows to represent documents and words in the same semantic space. The method can be applied with word embedding learned on the data with any state-of-the-art method [6, 12], or with embeddings that were previously learned¹ to reduce the computation cost. Contrary to many existing methods that make use of deep and complex neural networks (see Sect. 2 for related works), our method is fast, and it has only one parameter to tune.

We propose to construct a weight vector for each document using both textual and network information. We can then project the documents into the prelearned word vector space using this vector (see Fig. 1). The method is straightforward to apply, as it only requires applying well studied word embedding methods and matrix multiplication. We show in Sect. 4 that it outperforms or matches existing methods in classification and link prediction tasks and we demonstrate that projecting the documents into the word embedding space can provide semantic insights.




Fig. 1. Our method performs smoothing (represented as red arrows) on the documents' centroid representations (the square blocks). As the document in the blue circle (dots are words) is connected to the orange one, their representations get closer. The document in the green circle is isolated, thus it remains unchanged by the smoothing effect. (Color figure online)

2 Related Work

Several methods study the embedding of paragraph or short documents such as [10], generalizing the seminal word2vec models proposed by [12]. These approaches go beyond the simple method that consists in building a weighted average of representations of words that compose the document. For example in [2], authors propose to perturb weights for word average projection using Singular Value Decomposition (SVD). This last approach inspired our work as they show that word average is often a relevant baseline that can be improved in some cases using contextual smoothing.

As stated above, many corpora are structured in networks, providing additional information on documents semantics. TADW [20] is the first method that

¹ E.g., <https://fasttext.cc/>.

deals with this kind of data. It formulates network embedding [15] as a matrix tri-factorization problem to integrate textual information. Subsequent methods mainly adopt neural network based models: STNE [11] extends the seq2seq models, Graph2Gauss [3] learns both representations and variances via energy based learning, and VGAE [8] adopts a variational encoder. Even if these approaches yield good results, they require tuning a lot of hyperparameters. Two methods are based on factorization approaches: GVN-r-t [4], that extends GloVe [14], and AANE [7]. None of these methods learn documents and words embedding in the same space. In [10] and [1], authors represent them in a comparable space. Yet, they do not consider network information, as opposed to LDE [19]. Nonetheless, this last method requires labels associated with nodes, making it a supervised approach. Our method projects the documents and the words into the same space in an unsupervised fashion, with only one hyperparameter to tune. We will now present the formulation of this approach.

3 RLE: Document Projection with Smoothing

In this section, we present our model to build vector representations for a collection of linked documents. From now on, we will refer to our method as Regularized Linear Embedding (RLE). Matrices are in capital letters, and if X is a matrix, we write x_i the i -th row of X . From a network of n nodes, we extract a pairwise similarity matrix $S \in \mathbb{R}^{n \times n}$, computed as $S = \frac{A+A^2}{2}$ with A the transition matrix of the graph. Similarly to [20], this matrix considers both first and second order similarities. v is the number of words in the vocabulary. The corpus is represented as a document-term matrix $T \in \mathbb{R}^{n \times v}$, with each entry of T being the relative frequency of a word in a given document.

With $U \in \mathbb{R}^{v \times k}$ a matrix of pretrained word embeddings in dimension k , our goal is to build a matrix $D \in \mathbb{R}^{n \times k}$ of document embeddings, in the same space as the word embeddings. We build, for each document, a weight vector $p_i \in \mathbb{R}^v$, stacked in a matrix P and define the embedding of a document as $d_i = p_i U$. We construct p_i as follows: we first compute a smoothing matrix $B \in \mathbb{R}^{n \times v}$ with:

$$b_i = \frac{1}{\sum_j S_{i,j}} \sum_j S_{i,j} t_j. \quad (1)$$

Each row b_i of this matrix is a centroid of the initial document-term frequency matrix T , weighted by the similarity between the document i and each of the other documents. Then, we compute the weight matrix P according to T and B , in matrix notation:

$$P = (1 - \lambda)T + \lambda B, \quad (2)$$

where $\lambda \in [0, 1]$ controls the smoothing intensity. Then, we compute $D = PU$. Our method implies matrix multiplication and normalization only, making it fast and easily scalable. When $\lambda = 0$, $P = T$, thus, we recover the word average method. When $\lambda = 1$, we obtain $P = B$ and thus embed the documents with respect to the contextual information only (i.e., the similar documents). We illustrate the effect of smoothing in Fig. 1.

4 Experiments

In this section, we present our experimental results on classification and link prediction tasks, followed by a qualitative analysis of document representations.

We use two citation networks: Cora [18] and DBLP [13,17]. We also use New York Times articles (<https://www.nytimes.com/>) from January 2007. We create a link between pairs of articles sharing a common tag. The class corresponds to the article section. Cora contains 2,211 labeled documents (7 classes) with 5,001 citation links. The dataset includes the abstract of each article. The New York Times dataset (Nyt) contains 5,135 documents, 3,050,513 edges and 4 classes. Dblp has 60,744 documents (4 classes) and 52,914 edges. It includes the title of the articles only. After pruning the vocabulary (removing stop words, filtering word occurring in more than 25% of the corpus and less than 10 times), we obtain vocabularies made of 2,421 features for the Cora dataset, 6,407 for the Nyt dataset, and 3,763 for Dblp.

All embeddings are in dimension 160. We use DeepWalk with 40 walks of length 40, and a window of size 10. We also experiment with Latent Semantic Analysis (LSA) [5] and a concatenation of LSA and DeepWalk representations in dimension 80 as done by [20], referred as “Concatenation”. We also compare the performance of RLE with recent methods that embed attributed networks: STNE, Graph2Gauss, GVNR-t, VGAE, AANE and TADW. For STNE, we set the depth to 1 which leads to the best scores in our experiments. For Graph2Gauss, we set $K=1$, $\text{depth}=1$. We use default architecture for VGAE and determine optimal λ and ρ for AANE, and x_{min} for GVNR-t. For TADW, we use LSA in dimension 200 as a textual feature matrix and set regularization to 0.2, following authors’ recommendation. For each method, we use the implementation provided by the authors. We discard LDE since it is semi supervised and will not lead to a fair comparison.

RLE needs prelearned word representations. Hence, we build word vectors using Skip-gram with negative sampling [12]. We use the implementation in gensim², with window size of 15 for Cora, 10 for Nyt and 5 for DBLP, and 5 negative examples for both. The procedure is fast (46s for Cora, 84 on DBLP and 42 on Nyt). Similarly to baselines methods, we use the value of λ (0.7) that produces the optimal results on both datasets (see Fig. 2).




Fig. 2. Impact of λ on RLE in terms of document classification for $d = 160$. Optimum is achieved around 0.7 on each dataset (Cora, Nyt: 0.7, Dblp: 0.65).

² <https://radimrehurek.com/gensim/>.

4.1 Quantitative Results

We evaluate RLE in its ability to separate documents by classes in the embedding space and to predict links between documents. We perform SVM with L2 regularization on the vector representations of documents and report Micro F1 scores for different train/test ratios in Table 1. The regularisation strength is fixed through grid search. We also report computation times in second. For link prediction, we hide a percent of edges and compare the cosine similarity between hidden pairs and negative examples of unconnected documents. We report the Area Under the Roc Curve in Table 2.

Table 1. Comparison of Micro-F1 results on a classification task for different train/test ratios. The best score is in bold, second best is underlined. Execution time order is presented in seconds (Time).

train/test ratio	Cora				Dblp			
	10%	30%	50%	Time	10%	30%	50%	Time
DeepWalk	70.6 (2.0)	77.2 (0.9)	81.0 (0.7)	10 ¹	52.3 (0.4)	53.4 (0.1)	53.5 (0.2)	10 ²
LSA	72.3 (1.9)	79.0 (0.7)	80.6 (0.7)	10 ⁻²	73.5 (0.2)	74.1 (0.1)	74.2 (0.2)	10 ¹
Concatenation	71.4 (2.1)	80.5 (1.0)	84.0 (1.1)	10 ¹	<u>77.5</u> (0.2)	<u>78.0</u> (0.1)	<u>78.2</u> (0.2)	10 ²
TADW	81.9 (0.8)	86.3 (0.8)	87.4 (0.8)	10 ⁻¹	74.8 (0.1)	75.3 (0.2)	75.5 (0.1)	10 ¹
AANE	79.8 (0.9)	83.3 (1.1)	84.4 (0.7)	10 ⁻¹	73.3 (0.1)	73.9 (0.1)	74.2 (0.2)	10 ²
GVNR-t	<u>83.7</u> (1.2)	<u>86.4</u> (0.7)	87.0 (0.8)	10 ¹	69.6 (0.1)	70.1 (0.1)	70.2 (0.2)	10 ²
VGAE	72.3 (1.7)	79.2 (0.9)	81.1 (0.7)	10 ¹	Memory overflow			–
G2G	79.0 (1.5)	83.7 (0.8)	84.8 (0.7)	10 ¹	70.8 (0.1)	71.3 (0.2)	71.5 (0.2)	10 ²
STNE	79.4 (1.0)	84.7 (0.7)	86.7 (0.8)	10 ²	73.8 (0.2)	74.4 (0.1)	74.5 (0.1)	10 ⁴
RLE	84.0 (1.3)	86.9 (0.5)	87.7 (0.6)	10 ¹	79.8 (0.2)	80.9 (0.2)	81.2 (0.1)	10 ¹

train/test ratio	Nyt			
	10%	30%	50%	Time
DeepWalk	66.9 (0.7)	68.2 (0.3)	68.7 (0.9)	10 ²
LSA	71.6 (1.0)	75.7 (0.7)	76.7 (0.7)	10 ⁻²
Concatenation	77.9 (0.3)	80.0 (0.5)	81.1 (0.7)	10 ²
TADW	75.8 (0.5)	78.4 (0.5)	79.4 (0.4)	10 ¹
AANE	71.7 (0.5)	75.6 (0.8)	76.9 (1.1)	10 ¹
GVNR-t	74.3 (0.4)	76.0 (0.6)	76.7 (0.6)	10 ²
VGAE	68.1 (0.8)	69.3 (0.9)	70.1 (0.6)	10 ²
G2G	69.0 (0.5)	70.5 (0.7)	71.5 (0.8)	10 ²
STNE	75.1 (0.7)	77.3 (0.5)	78.1 (0.6)	10 ²
RLE	<u>77.7</u> (0.7)	<u>79.3</u> (0.5)	<u>80.0</u> (0.6)	10 ¹

In the classification task, RLE outperforms existing methods on Cora and Dblp, and is the second best method on Nyt. Interestingly, GVNR-t performs well with few training example, while TADW become second with 50% of training examples. Let us highlight that RLE runs fast, it is even faster than AANE on Dblp. Additionally, it is up to four orders of magnitude faster than STNE on Dblp. Additionally Fig. 2 shows that the optimal lambda values are similar for both datasets. Its tuning is not that crucial since RLE outperforms the baselines with $\lambda \in [0.6, 0.85]$ on Cora, $\lambda \in [0.15, 0.85]$ on DBLP, and every methods except Concatenation for $\lambda \in [0.45, 0.8]$ on Nyt.

Table 2. Comparison of AUC results on a link prediction task for different percents of edges hidden. The best score is in bold, second best is underlined.

% edges hidden	Cora		Dblp	
	50%	25%	50%	25%
DeepWalk	73.2 (0.6)	80.9 (1.0)	89.7 (0.0)	93.2 (0.2)
LSA	87.4 (0.6)	87.2 (0.8)	54.2 (0.1)	54.8 (0.0)
Concatenation	77.9 (0.3)	83.7 (0.8)	88.8 (0.0)	<u>92.6</u> (0.3)
TADW	90.1 (0.4)	93.3 (0.4)	61.2 (0.1)	65.0 (0.5)
AANE	83.1 (0.8)	86.6 (0.8)	67.4 (0.1)	66.5 (0.1)
GVNR-t	83.9 (0.9)	91.5 (1.1)	88.1 (0.3)	91.4 (0.1)
VGAE	87.1 (0.4)	88.2 (0.7)	Does not scale	
Graph2Gauss	<u>92.0</u> (0.3)	<u>93.8</u> (1.0)	88.0 (0.1)	92.1 (0.5)
STNE	83.1 (0.5)	90.0 (1.0)	45.6 (0.0)	53.4 (0.1)
RLE	94.3 (0.2)	94.8 (0.2)	<u>89.3</u> (0.1)	91.2 (0.2)

In link prediction, RLE outperforms existing methods on Cora, while DeepWalk yields better results than baselines on Dblp. This might be due to the shortness of the documents (mean length is 6 while it is 49 for Cora): the textual information may not be as informative as the network information for link prediction.

4.2 Qualitative Insights

We compute a vector representation for a class by computing the centroid of the representations of the documents inside this class. We present the closest words to this representation in term of cosine similarity, which provides a general description of the class. In Table 3, we present a description using this method for the first four classes of the Cora Dataset. We also provide most weighted terms when computing the mean of documents $tf \cdot idf$ of the class. The $tf \cdot idf$ method produces too general words, such as “learning”, “algorithm” and “model”. RLE seems to provide specific words, which makes the descriptions more relevant.

Table 3. Classes description with our method as opposed to $tf \cdot idf$. Words that are repeated across classes are in bold. RLE produces more discriminative descriptions

Cora							
Class 1		Class 2		Class 3		Class 4	
RLE	$tf \cdot idf$	RLE	$tf \cdot idf$	RLE	$tf \cdot idf$	RLE	$tf \cdot idf$
hebbian	neural	reinforcement	learning	posterior	bayesian	pac	learning
network	network	discounted	reinforcement	gibbs	model	learnability	algorithm
layers	networks	qlearning	control	bayesian	models	polynomialtime	algorithms
multilayer	learning	rl	state	mcmc	algorithm	dnf	model
filters	model	multiagent	policy	sampler	belief	queries	decision

5 Conclusion

In this article, we presented the RLE method for embedding documents that are organized in a network. Despite its simplicity, RLE shows state-of-the-art results for the three considered datasets. It is faster than most recent deep-learning methods. Furthermore, it provides informative qualitative insights. Future works will concentrate on automatically tuning λ , and exploring the effect of the similarity matrix S .

References

1. Ailem, M., Salah, A., Nadif, M.: Non-negative matrix factorization meets word embedding. In: Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 1081–1084. ACM (2017)
2. Arora, S., Liang, Y., Ma, T.: A simple but tough-to-beat baseline for sentence embeddings. In: International Conference on Learning Representations (2016)
3. Bojchevski, A., Günnemann, S.: Deep Gaussian embedding of graphs: unsupervised inductive learning via ranking. In: Proceeding of the International Conference on Learning Representations. ICLR (2018)
4. Brochier, R., Guille, A., Velcin, J.: Global vectors for node representations. In: Proceedings of the World Wide Web Conference, pp. 2587–2593. WWW (2019)
5. Deerwester, S., Dumais, S.T., Furnas, G.W., Landauer, T.K., Harshman, R.: Indexing by latent semantic analysis. *J. Am. Soc. Inf. Sci.* **41**(6), 391–407 (1990)
6. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: BERT: pre-training of deep bidirectional transformers for language understanding. arXiv preprint [arXiv:1810.04805](https://arxiv.org/abs/1810.04805) (2018)
7. Huang, X., Li, J., Hu, X.: Accelerated attributed network embedding. In: Proceedings of the SIAM International Conference on Data Mining, pp. 633–641. SDM (2017)
8. Kipf, T.N., Welling, M.: Variational graph auto-encoders. In: Bayesian Deep Learning Workshop, BDL-NeurIPS (2016)
9. Kuzi, S., Shtok, A., Kurland, O.: Query expansion using word embeddings. In: Proceedings of the 25th ACM International Conference on Information and Knowledge Management, pp. 1929–1932. ACM (2016)
10. Le, Q., Mikolov, T.: Distributed representations of sentences and documents. In: International Conference on Machine Learning, pp. 1188–1196 (2014)
11. Liu, J., He, Z., Wei, L., Huang, Y.: Content to node: self-translation network embedding. In: Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 1794–1802. ACM (2018)
12. Mikolov, T., Sutskever, I., Chen, K., Corrado, G.S., Dean, J.: Distributed representations of words and phrases and their compositionality. In: Advances in Neural Information Processing Systems, pp. 3111–3119 (2013)
13. Pan, S., Wu, J., Zhu, X., Zhang, C., Wang, Y.: Tri-party deep network representation. In: Proceedings of the International Joint Conference on Artificial Intelligence, pp. 1895–1901. IJCAI (2016)
14. Pennington, J., Socher, R., Manning, C.: GloVe: global vectors for word representation. In: Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP), pp. 1532–1543 (2014)

15. Perozzi, B., Al-Rfou, R., Skiena, S.: DeepWalk: online learning of social representations. In: Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 701–710. ACM (2014)
16. Seyler, D., Chandar, P., Davis, M.: An information retrieval framework for contextual suggestion based on heterogeneous information network embeddings. In: The 41st International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 953–956. ACM (2018)
17. Tang, J., Zhang, J., Yao, L., Li, J., Zhang, L., Su, Z.: ArnetMiner: extraction and mining of academic social networks. In: Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 990–998. KDD (2008)
18. Tu, C., Liu, H., Liu, Z., Sun, M.: CANE: context-aware network embedding for relation modeling. In: Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), vol. 1, pp. 1722–1731 (2017)
19. Wang, S., Tang, J., Aggarwal, C., Liu, H.: Linked document embedding for classification. In: Proceedings of the 25th ACM International on Conference on Information and Knowledge Management, pp. 115–124. ACM (2016)
20. Yang, C., Liu, Z., Zhao, D., Sun, M., Chang, E.Y.: Network representation learning with rich text information. In: International Joint Conference on Artificial Intelligence (2015)