

Context encoders as a simple but powerful extension of word2vec

Franziska Horn

Machine Learning Group

Technische Universität Berlin, Germany

franziska.horn@campus.tu-berlin.de

Abstract

With a strikingly simple architecture and the ability to learn meaningful word embeddings efficiently from texts containing billions of words, word2vec remains one of the most popular neural language models used today. However, as only a single embedding is learned for every word in the vocabulary, the model fails to optimally represent words with multiple meanings and, additionally, it is not possible to create embeddings for new (out-of-vocabulary) words on the spot. Based on an intuitive interpretation of the continuous bag-of-words (CBOW) word2vec model’s negative sampling training objective in terms of predicting context based similarities, we motivate an extension of the model we call context encoders (ConEc). By multiplying the matrix of trained word2vec embeddings with a word’s average context vector, out-of-vocabulary (OOV) embeddings and representations for words with multiple meanings can be created based on the words’ local contexts. The benefits of this approach are illustrated by using these word embeddings as features in the CoNLL 2003 named entity recognition (NER) task.

1 Introduction

Representation learning is very prominent in the field of natural language processing (NLP). For example, word embeddings learned by neural language models (NLM) were shown to improve the performance when used as features for supervised learning tasks such as named entity recognition (NER) (Collobert et al., 2011; Turian et al., 2010). The popular *word2vec* model (Mikolov et al., 2013a,b) learns meaningful word embed-

dings by considering only the words’ local contexts and thanks to its shallow architecture it can be trained very efficiently on large corpora. The model, however, only learns a single representation for words from a fixed vocabulary. This means, if in a task we encounter a new word that was not present in the texts used for training, we cannot create an embedding for this word without repeating the time consuming training procedure of the model.¹ Additionally, a single embedding does not optimally represent words with multiple meanings. For example, “Washington” is both the name of a US state as well as a former president and only by taking into account the word’s local context one can identify the proper sense.

Based on an intuitive interpretation of the continuous bag-of-words (CBOW) word2vec model’s negative sampling training objective, we propose an extension of the model we call *context encoders* (ConEc). This allows for an easy creation of OOV embeddings as well as a better representation of words with multiple meanings simply by multiplying the trained word2vec embeddings with the words’ average context vectors. As demonstrated on the CoNLL 2003 NER challenge, using the word embeddings created with ConEc instead of word2vec as features improves the classification performance significantly.

Related work In the past, NLM have addressed the issue of polysemy in various ways. For example, *sense2vec* is an extension of word2vec, where in a preprocessing step all words in the training corpus are annotated with their part-of-speech (POS)

¹In practice the model is trained on such a large vocabulary that it is rare to encounter a word that does not have an embedding. Yet there are still scenarios where this is the case, for example, it is unlikely that the term “W10281545” is encountered in a regular training corpus, but we might still want its embedding to represent a search query like “whirlpool W10281545 ice maker part”.

tag and then the embeddings are learned for tokens consisting of the words themselves and their POS tags, thereby generating different representations e.g. for words that are used both as a noun and verb (Trask et al., 2015). Other methods first cluster the contexts the words appear (Huang et al., 2012) or use additional resources such as wordnet to identify multiple meanings of words (Rothe and Schütze, 2015). One possibility to create OOV embeddings is to learn representations for all character n-grams in the texts and then compute the embedding of a word by combining the embeddings of the n-grams occurring in it (Bojanowski et al., 2016). However, none of these NLM are designed to solve both the OOV and polysemy problem at the same time and compared to word2vec they require more parameters, resources, or additional steps in the training procedure. ConEc on the other hand can generate OOV embeddings as well as better representations for words with multiple meanings simply by multiplying the matrix of trained word2vec embeddings with the words’ average context vectors.

2 Background: CBOW word2vec trained with negative sampling

Word2vec learns d -dimensional vector representations, referred to as word embeddings, for all N words in the vocabulary. It is a shallow NLM with parameter matrices $W_0, W_1 \in \mathbb{R}^{N \times d}$, which are tuned iteratively by scanning huge amounts of texts sentence by sentence. Based on some context words the algorithm tries to predict the target word between them. Mathematically, this is realized by first computing the sum of the embeddings of the context words by selecting the appropriate rows from W_0 . This vector is then multiplied by several rows selected from W_1 : one of these rows corresponds to the target word, while the others correspond to k ‘noise’ words, selected at random (negative sampling). After applying a non-linear activation function, the backpropagation error is computed by comparing this output to a label vector $\mathbf{t} \in \mathbb{R}^{k+1}$, which is 1 at the position of the target word and 0 for all k noise words. After the training of the model is complete, the word embedding for a target word is the corresponding row of W_0 .

3 Context Encoders

Similar words appear in similar contexts (Harris, 1954), for example, two words synonymous with

each other could be exchanged for one another in almost all contexts without a reader noticing. Based on the context word co-occurrences, pairwise similarities between all N words of the vocabulary can be computed, resulting in a similarity matrix $S \in \mathbb{R}^{N \times N}$ (or for a single word w the vector $\mathbf{s}_w \in \mathbb{R}^N$) with similarity scores between 0 and 1. These similarities should be preserved in the word embeddings, e.g. the cosine similarity between the embedding vectors of two words used in similar contexts should be close to 1, or, more generally, the scalar product of the matrix with word embeddings $Y \in \mathbb{R}^{N \times d}$ should approximate S . Of course, the most straightforward way of obtaining word embeddings satisfying $YY^T \approx S$ would be to compute the singular value decomposition (SVD) of the similarity matrix S and use the eigenvectors corresponding to the d largest eigenvalues (Levy et al., 2014, 2015). As our vocabulary typically comprises several 10,000 words, however, performing an SVD of the corresponding similarity matrix is computationally far too expensive. Yet, while the similarity matrix would be huge, it would also be quite sparse, as many words are of course not synonymous with each other. If we picked a small number k of random words, chances are their similarities to a target word would be close to 0. So, while the product of a single word’s embedding $\mathbf{y}_w \in \mathbb{R}^d$ and the matrix of all embeddings Y should result in a vector $\hat{\mathbf{s}}_w \in \mathbb{R}^N$ close to the true similarities \mathbf{s}_w of this word, if we only consider a small subset of $\hat{\mathbf{s}}_w$ corresponding to the word itself and k random words, it is sufficient if this approximates the binary vector $\mathbf{t}_w \in \mathbb{R}^{k+1}$, which is 1 for the word itself and 0 elsewhere.

The CBOW word2vec model trained with negative sampling can therefore be interpreted as a neural network (NN) that predicts a word’s similarities to other words (Fig. 1). During training, for each occurrence i of a word w in the texts, a binary vector $\mathbf{x}_{w_i} \in \mathbb{R}^N$, which is 1 at the positions of the context words of w and 0 elsewhere, is used as input to the network and multiplied by a set of weights W_0 to arrive at an embedding $\mathbf{y}_{w_i} \in \mathbb{R}^d$ (the summed rows of W_0 corresponding to the context words). This embedding is then multiplied by another set of weights W_1 , which corresponds to the full matrix of word embeddings Y , to produce the output of the network, a vector $\hat{\mathbf{s}}_{w_i} \in \mathbb{R}^N$ containing the approximated similarities of the word w to all other words. The training error is then

computed by comparing a subset of the output to a binary target vector $\mathbf{t}_{w_i} \in \mathbb{R}^{k+1}$, which serves as an approximation of the true similarities \mathbf{s}_w when considering only a small number of random words. We refer to this interpretation of the model as *context encoders* (ConEc), as it is closely related to similarity encoders (SimEc), a dimensionality reduction method used for learning similarity preserving representations of data points (Horn and Müller, 2017).




Figure 1: Context encoder (ConEc) NN architecture corresponding to the CBOW word2vec model trained with negative sampling.

While the training procedure of ConEc is identical to that of word2vec, there is a difference in the computation of a word’s embedding after the training is complete. In the case of word2vec, the word embedding is simply the row of the tuned W_0 matrix. When considering the idea behind the optimization procedure, however, we instead propose to create the representation of a target word w by multiplying W_0 with the word’s average context vector \mathbf{x}_w , as this better resembles how the word embeddings are computed during training.

We distinguish between a word’s ‘global’ and ‘local’ average context vector (CV): The global CV is computed as the average of all binary CVs \mathbf{x}_{w_i} corresponding to the M_w occurrences of w in the whole training corpus:

$$\mathbf{x}_{w_{\text{global}}} = \frac{1}{M_w} \sum_{i=1}^{M_w} \mathbf{x}_{w_i},$$

while the local CV $\mathbf{x}_{w_{\text{local}}}$ is computed likewise but considering only the m_w occurrences of w in a single document. We can now compute the embedding of a word w by multiplying W_0 with the

weighted average between both CVs:

$$\mathbf{y}_w = (a \cdot \mathbf{x}_{w_{\text{global}}} + (1 - a) \mathbf{x}_{w_{\text{local}}})^\top W_0 \quad (1)$$

with $a \in [0, 1]$. The choice of a determines how much emphasis is placed on the word’s local context, which helps to distinguish between multiple meanings of the word (Melamud et al., 2015).² As an out-of-vocabulary word does not have a global CV (as it never occurred in the training corpus), its embedding is computed solely based on the local context, i.e. setting $a = 0$.

With this new perspective on the model and optimization procedure, another advancement is feasible. Since the context words are merely a sparse feature vector used as input to a NN, there is no reason why this input vector should not contain other features about the target word as well. For example, the feature vector \mathbf{x}_w could be extended to contain information about the word’s case, part-of-speech (POS) tag, or other relevant details. While this would increase the dimensionality of the first weight matrix W_0 to include the additional features when mapping the input to the word’s embedding, the training objective and therefore also W_1 would remain unchanged. These additional features could be especially helpful if details about the words would otherwise get lost in preprocessing (e.g. by lowercasing) or to retain information about a word’s position in the sentence, which is ignored in a BOW approach. These extended ConEcs are expected to create embeddings that distinguish even better between the words’ different senses by taking into account, for example, if the word is used as a noun or verb in the current context, similar to the sense2vec algorithm (Trask et al., 2015). But instead of learning multiple embeddings per term explicitly, like sense2vec, only the dimensionality of the input vector is increased to include the POS tag of the current word as a feature, which is expected to improve generalization if few training examples are available.

4 Experiments

The word embeddings learned by word2vec and context encoders are evaluated on the CoNLL 2003 NER benchmark task (Tjong et al., 2003). We use a CBOW word2vec model trained with negative sampling as described above where $k = 13$, the embedding dimensionality d is 200 and we use a context window of 5 words. The word embeddings

²This implicitly assumes a word is only used in a single sense in one document.




Figure 2: Results of the NER task based on three random initializations of the word2vec model. *Left panel:* Overall results, where the mean performance using word2vec embeddings (*dashed lines*) is considered as our baseline, all other embeddings are computed with ConEcs using various combinations of the words’ global and local CVs. *Right panel:* Increased performance (mean and standard deviation) on the test fold when using ConEc: Multiplying the word2vec embeddings with global CVs yields a performance gain of 2.5 percentage points (A). By additionally using local CVs to create OOV word embeddings we gain another 1.7 points (B). When using a combination of global and local CVs (with $a = 0.6$) to distinguish between the different meanings of words, the F1-score increases by another 5.1 points (C), yielding a F1-score of 39.92%, which marks a significant improvement compared to the 30.59% reached with word2vec features.

created by ConEc are built directly on top of the word2vec model by multiplying the original embeddings (W_0) with the respective context vectors. Code to replicate the experiments is available online.³

Named Entity Recognition The main advantage of context encoders is that they can use local context to create OOV embeddings and distinguish between the different senses of words. The effects of this are most prominent in a task such as NER, where the local context of a word can make all the difference, e.g. to distinguish between the “Chicago Bears” (an organization) and the city of Chicago (a location). We tested this on the CoNLL 2003 NER task by using the word embeddings as features together with a logistic regression classifier. The reported F1-scores were computed using the official evaluation script. The results achieved with various word embeddings on the training, development and test part of the CoNLL task are reported in Fig. 2. Please note that we are using this task as an extrinsic evaluation to illustrate the advantages of ConEc embeddings over the regular word2vec embeddings. To isolate the effects on the performance, we are only using these word embeddings

as features, while of course the performance on this NER challenge is typically much higher when other features such as a word’s case or POS tag are included as well.

The word2vec embeddings were trained on the documents used in the training part of the task and OOV words in the development and test parts are represented as zero vectors.⁴ With three parameter settings we illustrate the advantages of ConEc: A) *Multiplying the word2vec embeddings by the words’ average context vectors generally improves the embeddings.* To show this, ConEc word embeddings were computed using only global CVs (Eq. 1 with $a = 1$), which means OOV words again have a zero representation. With these embeddings (labeled ‘global’ in Fig. 2) the performance improves on the dev and test folds of the task.

B) *Useful OOV embeddings can be created from the local context of a new word.* To show this, the ConEc embeddings for words from the training vocabulary ($w \in N$) were computed as in A), but now the embeddings for OOV words ($w' \notin N$) were computed using local CVs (Eq. 1 with $a = 1 \forall w \in N$ and $a = 0 \forall w' \notin N$; referred to as ‘OOV’ in the figure). The training performance

³<https://github.com/cod3licious/conec>

⁴Since this is a very small corpus, we trained word2vec for 25 iterations on these documents.

stays the same, of course, as here all words have an embedding based on their global contexts, but there is a jump in the ConEc performance on the dev and test folds, where OOV words now have a representation based on their local contexts.

C) *Better embeddings for a word with multiple meanings can be created by using a combination of the word’s average global and local CVs as input to the ConEc.* To show this, the OOV embeddings were computed as in B), but now for the words occurring in the training vocabulary, the local context was taken into account as well by setting $a < 1$ (Eq. 1 with $a \in [0, 1) \forall w \in N$ and $a = 0 \forall w' \notin N$). The best performances on all folds are achieved when averaging the global and local CVs with around $a = 0.6$ before multiplying them with the word2vec embeddings, which clearly shows that ConEc embeddings created by incorporating local context can help distinguish between multiple meanings of words.

5 Conclusion

Context encoders are a simple but powerful extension of the CBOW word2vec model trained with negative sampling. By multiplying the matrix of trained word2vec embeddings with the words’ average context vectors, ConEcs are able to easily create OOV embeddings on the spot as well as distinguish between multiple meanings of words based on their local contexts. The benefits of this were demonstrated on the CoNLL NER challenge.

Acknowledgments

I would like to thank Antje Relitz, Ivana Balažević, Christoph Hartmann, Klaus-Robert Müller, and other anonymous reviewers for their helpful comments on earlier versions of this manuscript. Franziska Horn acknowledges funding from the Elsa-Neumann scholarship from the TU Berlin.

References

Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2016. Enriching word vectors with subword information. *arXiv preprint arXiv:1607.04606*.

Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011. Natural language processing (almost) from scratch. *The Journal of Machine Learning Research* 12:2493–2537.

Zellig S Harris. 1954. Distributional structure. *Word* 10(2-3):146–162.

Franziska Horn and Klaus-Robert Müller. 2017. Learning similarity preserving representations with neural similarity encoders. *arXiv preprint arXiv:1702.01824*.

Eric H Huang, Richard Socher, Christopher D Manning, and Andrew Y Ng. 2012. Improving word representations via global context and multiple word prototypes. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Long Papers-Volume 1*. ACL, pages 873–882.

Omer Levy, Yoav Goldberg, and Ido Dagan. 2015. Improving distributional similarity with lessons learned from word embeddings. *Transactions of the Association for Computational Linguistics* 3:211–225.

Omer Levy, Yoav Goldberg, and Israel Ramat-Gan. 2014. Linguistic regularities in sparse and explicit word representations. In *CoNLL*. pages 171–180.

Oren Melamud, Ido Dagan, and Jacob Goldberger. 2015. Modeling word meaning in context with substitute vectors. In *Human Language Technologies: The 2015 Annual Conference of the North American Chapter of the ACL*.

Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013a. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.

Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013b. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*. pages 3111–3119.

Sascha Rothe and Hinrich Schütze. 2015. Autoextend: Extending word embeddings to embeddings for synsets and lexemes. *arXiv preprint arXiv:1507.01127*.

EF Tjong, Kim Sang, and F De Meulder. 2003. Introduction to the CoNLL-2003 Shared Task: Language-Independent Named Entity Recognition. In Walter Daelemans and Miles Osborne, editors, *Proceedings of CoNLL-2003*. Edmonton, Canada, pages 142–147.

Andrew Trask, Phil Michalak, and John Liu. 2015. sense2vec-a fast and accurate method for word sense disambiguation in neural word embeddings. *arXiv preprint arXiv:1511.06388*.

Joseph Turian, Lev Ratinov, and Yoshua Bengio. 2010. Word representations: a simple and general method for semi-supervised learning. In *Proceedings of the 48th annual meeting of the association for computational linguistics*. Association for Computational Linguistics, pages 384–394.