

Using dependency parsing for few-shot learning in distributional semantics

Stefania Preda

University College London
United Kingdom

stefipredacs@gmail.com

Guy Emerson

University of Cambridge
United Kingdom

gete2@cam.ac.uk

Abstract

In this work, we explore the novel idea of employing dependency parsing information in the context of few-shot learning, the task of learning the meaning of a rare word based on a limited amount of context sentences. Firstly, we use dependency-based word embedding models as background spaces for few-shot learning. Secondly, we introduce two few-shot learning methods which enhance the additive baseline model by using dependencies.

1 Introduction

Distributional semantics models create word embeddings based on the assumption that the meaning of a word is defined by the contexts it is used in (for an overview, see: [Sahlgren, 2008](#); [Lenci, 2018](#); [Boleda, 2020](#); [Emerson, 2020](#)). A fundamental challenge for these approaches is the difficulty of producing high-quality embeddings for rare words, since the models often require vast amounts of training examples ([Adams et al., 2017](#); [Van Haute et al., 2019](#)). To address this problem, various few-shot learning methods have been previously introduced. The goal of a few-shot learning technique is to learn an embedding that captures the meaning of a word, given only a few context sentences. The rare word’s vector has to be placed in an existing *background* space of embeddings.

Few-shot learning in distributional semantics is a relatively underexplored area, with important practical applications. Having good representations of rare words is highly desirable in applications aiming to understand dialects or regionalisms, as well as specific technical language.

In this work, we explore the idea of incorporating information from the dependency parse of sentences in the context of few shot-learning. An intuition why this might be useful is provided in [Figure 1](#). In the given sentence, the most relevant word for inferring the meaning of the target rare word “conflagration” is “destroyed”. Even if this

word is located far from the target, it is directly connected to it through a nominal subject dependency. Moreover, the fact that the target word is used in a certain dependency structure might reveal important characteristics related to its meaning. Since in the case of few-shot learning the data is limited, using dependency parsing information is a resource with great potential to boost existing models.

As a first effort in this direction, this work provides three contributions. Firstly, we explore the effect of using dependency-based word embeddings as background spaces. Secondly, we introduce new few-shot learning methods leveraging the dependency parsing information. Lastly, we update a previous dependency-based background model to make it more suitable for few-shot learning.

2 Background: dependency-based word embeddings

The widely-used Skip-Gram model introduced by [Mikolov et al. \(2013\)](#) takes the contexts of a word to be those words surrounding it in a pre-defined window size. The model learns the embeddings in an unsupervised manner, using a feed-forward neural network trained on large amounts of sentences.

[Levy and Goldberg \(2014\)](#) proposed a different way to construct the contexts of a target word in the training process of the Skip-Gram model. Instead of taking the words from a pre-defined window, one takes the words that are connected to the target word by a syntactic dependency. The contexts were defined as the concatenation of the connected word and the label of the dependency. This allowed the model to differentiate between same words used in different syntactic roles.

The dependency-based word embeddings were found to be better at capturing similarity, while the window-based models capture relatedness. For example, a dependency-based model would produce close embeddings for “Rome” and “Florence”, which are *syntactically similar* since they can be




Figure 1: A dependency parse, illustrating that context words connected by a dependency can be important for inferring meaning in a few-shot setting, such as “destroyed” for the rare word “conflagration”.

used in the same grammatical contexts, while a window-based model is likely to place closely the embedding of highly related terms such as “Rome” and “ancient”, even if they cannot be used interchangeably since they are different parts of speech.

Levy and Goldberg’s model successfully captured syntactic similarity, but failed to express how different dependency types affect relations between words. Moreover, it introduced sparsity issues. Czarnowska et al. (2019) developed the Dependency Matrix model to address these shortcomings. Instead of incorporating the dependency labels in the context vocabulary, each dependency type d is associated with a matrix T_d , which acts as a meaning representation of the link between the target and the context words. The matrices T_d , as well as the vectors holding the target vectors e and context vectors o , are learned during training. Let \mathcal{D} be the set of training examples given by tuples of target word t , context word c and dependency type d . For each tuple, we generate a set \mathcal{D}' of negative samples (t, c', d) by drawing context words c' from a noise distribution and maintaining the same target word t and dependency type d . The learning goal is to maximise the function in (1), where σ is the sigmoid function and e_t and o_c are the vectors of the target and context word.

$$\sum_{(t,c,d) \in \mathcal{D}} \left(\log \sigma(u^{t,c,d}) + \sum_{(t,c',d) \in \mathcal{D}'} \log \sigma(-u^{t,c',d}) \right) \quad (1)$$

$$u^{t,c,d} = e_t \cdot T_d \cdot o_c \quad (2)$$

3 Background: few-shot learning

As a straight-forward yet successful baseline, the vector of the rare word is estimated by the sum of the vectors of the words in contexts, as proposed by Lazaridou et al. (2017) and Herbelot and Baroni (2017). The latter noticed that not including the stop-words greatly improves the performance on the evaluation tasks. To optimise the performance of the additive model, Van Haute et al. (2019) proposed weighting the context words according to distance and frequency, as well as subtracting a “negative sampling” vector. These modifications

take hyperparameters that are important for Skip-Gram’s strong performance, such as number of negative samples k and window size n (Levy et al., 2015), and apply them to the few-shot setting. For each word w in the vocabulary V , with frequency $f(w)$ and distance m from the target rare word t , and for a frequency threshold τ , we calculate the subsampling weight $s(w)$, the window weight $r(w)$ and negative sampling coefficient $n(w)$.

$$s(w) = \min \left(1, \sqrt{\frac{\tau}{f(w)}} \right) \quad (3)$$

$$r(w) = \max \left(0, \frac{n - m + 1}{n} \right) \quad (4)$$

$$n(w) = \frac{f(w)^{0.75}}{\sum_{w \in V} f(w)^{0.75}} \quad (5)$$

Assume \mathcal{C} is the collection of non-stop context words for the given target rare word t and v_c is the vector in the background space for each $c \in \mathcal{C}$. The vector of the target rare word t will is:

$$v_t = \sum_{c \in \mathcal{C}} v_c^{\text{add}} \quad \text{where} \quad (6)$$

$$v_c^{\text{add}} = s(c)r(c) \left(v_c - k \sum_{w \in V} n(w)v_w \right) \quad (7)$$

More involved models have been proposed for the task of few-shot learning. Khodak et al. (2018) introduced A La Carte, which applies a linear transformation to the sum of the context words obtained by the additive model. The weights of the linear transformation are optimised based on the co-occurrence matrix of the corpus. Van Haute et al. (2019) takes this approach further in the Neural A La Carte model, by using a neural network with a hidden layer to produce a non-linear transformation matrix, which adds flexibility.

The meaning of a rare word can often be deduced from the word form itself. This information has been leveraged in few-shot learning models. For example, the Form-Context Model (Schick and Schütze, 2019) is a hybrid method which retrieves the weighted sum between the surface form embedding of the rare word, obtained using FastText

(Bojanowski et al., 2017) and the context-based embedding, produced using the A La Carte model.

In this paper, we focus on additive methods, which do not require additional training on few-shot learning examples. This keeps the inference fast and in line with the *true few-shot learning* setting proposed by Perez et al. (2021).

4 Dependency-based FSL methods

Dependency relations proved to be an informative tool in the context of creating distributional semantics models. Based on this success, we introduce two dependency-based few-shot learning methods which build on top of the Additive model. In this section, we assume we have already trained a background space of embeddings v_i for each word i . In our setup, we chose to consider only the target embeddings learnt by the aforementioned background models, i.e. $v_i = e_i$. Alternatively, one could use the concatenation of the target and context embeddings.

Dependency Additive Model The starting point of our methods is the assumption that the closer a word is to the target word in the dependency graph, the more relevant it is for inferring the target’s meaning, as seen in Figure 1.

Our method assigns weights for each word in the sentence by considering the distances from the rare word in the dependency parse. For each context word c , let d_c be the number of dependency links from the target rare word t to c in the parse. Note that we consider links in both directions. The inferred vector v_t of the rare word is the weighted sum of the vectors of context words, where the weight w_c of each context word c is given in (8). The weight is chosen so that it is inversely proportional to the distance from the target, and we add 1 in order to avoid discarding context words which are far from the target in the dependency tree.

$$v_t = \sum_{c \in C} w_c v_c^{\text{add}} \quad \text{where } w_c = 1 + \frac{1}{d_c} \quad (8)$$

Initially, we experimented with simply applying the coefficients w_c on the vectors of the context words v_c . However, a better performance was achieved when we incorporated the the weighting steps in (7), so we used v_c^{add} instead of v_c .

Dependency Matrix Additive Model The Dependency Additive model above does not take into

account the type of dependency on each edge in the graph, which, as we have seen, plays an important role in capturing the meaning of the words in relation to each other. We thus devised a strategy to make use of this information.

Czarnowska et al. proposed the idea of using the learnt dependency matrices of the Dependency Matrix model for the task of semantic composition, by multiplying word embeddings with matrices over chains of dependencies. We apply the same idea in the context of few-shot learning. More precisely, instead of giving a weight for each vector of a context word, we multiply it with corresponding dependency matrices on the chain of dependencies from the target to the context. To be able to do this based on the original Dependency Matrix model, we would have to take into account that when we advance in the dependency parse, we have to switch between using the context vector (retrieved from o) and target vector (retrieved from e).

To simplify this process, we modified the Dependency Matrix model to use only one embedding per word, instead of separate context and target embeddings.¹ This also reduces the training time of the model. More precisely, we have the same training loss as in (1), but (2) is replaced by:

$$u^{t,c,d} = v_t \cdot T_d \cdot v_c \quad (9)$$

Having trained this model, we then make use of the matrices T_d , optimised for each dependency type d . For the target rare word t and each non-stop context word c , Let $D(t, c)$ be the path of dependency types from t to c . The vector of the target rare word is calculated as:

$$v_t = \sum_{c \in C} v'_c \quad \text{where } v'_c = \left(\prod_{d \in D(t,c)} T_d \right) v_c^{\text{add}} \quad (10)$$

5 Experiments

In our setup, we considered three background models: window-based Skip-Gram, dependency-based Skip-Gram and the modified Dependency Matrix model which only uses one embedding for each

¹This cannot be applied to Skip-Gram without causing every word to predict itself as a context. To allow Skip-Gram to use only one vector per word, Zornin and Elistratova (2019) propose using an indefinite inner product, which corresponds to T in (9) being a diagonal matrix of 1s and -1 s. In a similar vein, Bertolini et al. (2021) propose a more radical simplification of the Dependency Matrix model, which uses matrices that are non-zero only on the diagonal and off-diagonal.

Backgr. Model	FSL Model	DN		Chimera			CRW				
		MRR	MR	L2	L3	L6	1	2	4	8	16
Skip-Gram	Additive	0.010	5312	0.12	0.19	0.20	0.11	0.12	0.13	0.15	0.15
	Dep. Additive	0.021	4007	0.13	0.20	0.21	0.12	0.13	0.14	0.15	0.16
Dependency Skip-Gram	Additive	0.023	4671	0.14	0.21	0.21	0.11	0.14	0.15	0.16	0.17
	Dep. Additive	0.027	3785	0.16	0.21	0.23	0.12	0.15	0.16	0.17	0.18
Dependency Matrix	Additive	0.017	3367	0.13	0.23	0.25	0.15	0.17	0.20	0.22	0.22
	Dep. Additive	0.034	3140	0.14	0.26	0.29	0.18	0.20	0.22	0.24	0.25
	DM Additive	0.019	3163	0.15	0.24	0.31	0.16	0.20	0.20	0.21	0.22

Table 1: Results for different combinations of background and few-shot learning model, on three evaluation datasets. The best result for each column is marked in bold. Higher is better for all columns except MR.

word. To allow a direct comparison, we trained them all on the WikiWoods (Flickinger et al., 2010) snapshot of English Wikipedia. The same hyperparameters were used: a dimensionality of 100, 15 negative samples, a batch size of 5, and an Adagrad optimiser with an initial learning rate of 0.025. For the dependency models, we used the universal dependency parser provided by spaCy (Honnibal et al., 2020). We applied the two few-shot methods we devised, as well as the Additive model with window weighting, subsampling and negative sampling described in §3. The hyperparameters were $t = 10^{-6}$, $k = 15$ and $n = 5$.

5.1 Few-shot learning tasks

Definitional Nonce (DN) This task (Herbelot and Baroni, 2017) provides a single definition sentence for each test word. The test words are frequent words, which have gold vectors of high quality in the background space. At evaluation time, a new vector is computed for each test word, based on the few-shot learning model. The rank of the gold vector relative to the inferred vector is then calculated, i.e. the number of words from the vocabulary which are closer to the inferred vector than the gold vector is. The distance metric is cosine similarity - the bigger the similarity, the smaller the distance. The metrics retrieved are the Mean Reciprocal Rank (MRR) and median rank.

Chimera The Chimera task (Lazaridou et al., 2017) provides non-existing words (chimeras) with 6 context sentences, as well as similarity scores between the chimera and other existing words. The way in which the dataset was built simulates few-shot learning for humans, since the participants of the experiment needed to infer the meaning of a word they never saw before and rate its similarity with other concepts, based only on the 6 context

sentences. Trials with 2, 4 and 6 context sentences are conducted. After each trial, the similarity scores between the inferred vector and the vectors of the words provided is compared against the human similarity scores by retrieving the Spearman’s rank correlation coefficient.

Contextual Rare Words (CRW) Like Chimera, the CRW task (Khodak et al., 2018) is based on human ratings between pairs of words. This time the pairs contain a rare word and a frequent one, with an assumed reliable embedding in the background model. For each rare word, 255 context sentences are provided. The vector is generated using the few-shot model for 1, 2, 4, 8, 16 context sentences, selected at random. For each such experiment, the similarity scores between the few-shot vector and the background embedding of the non-rare word are calculated and compared against the human scores using the Spearman’s rank correlation coefficient. The scores are averaged out across 10 random selections of context sentences.

5.2 Results and Discussion

The results in Table 1 show that the dependency-based background models performed better than window-based Skip-Gram on all three evaluation tasks. For all background models, applying the Dependency Additive technique consistently improved the results of the Additive model. For the DN task and DM background model, there were three cases where the Additive model gave a rank of over 30,000, while the Dependency Additive model gave a rank of 1 or 2, showing the method’s potential for sentences of specific structures. The DM additive model showed a promising result for the Chimera task, but was still outperformed by the Dependency Additive model, and its scores had the biggest variance across all combinations. This

suggests that more careful weighting might be required.

6 Conclusion

We investigated the use of dependency information for few-shot learning in distributional semantics. We found that dependency-based contexts are more useful than window-based contexts, with better performance across three evaluation datasets. We proposed a simplified version of the Dependency Matrix model, using only one vector per word, which makes it easier to apply in a few-shot setting.

An important next step would be to investigate the use of the proposed methods for other languages, since our work was limited to English data and it is possible that the dependency structure is more relevant for few-shot learning in the case of specific languages. In order to do such an analysis, one would additionally need to create test data for the few shot-learning tasks, which would require the participation of speakers of the selected languages.

In future work, performance might be further improved by training an A La Carte model (discussed in §3), where the use of dependencies would make it possible to use a graph-convolutional network (Marcheggiani and Titov, 2017).

References

- Oliver Adams, Adam Makarucha, Graham Neubig, Steven Bird, and Trevor Cohn. 2017. [Cross-lingual word embeddings for low-resource language modeling](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 937–947, Valencia, Spain. Association for Computational Linguistics.
- Lorenzo Bertolini, Julie Weeds, David Weir, and Qiwei Peng. 2021. [Representing syntax and composition with geometric transformations](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 3343–3353, Online. Association for Computational Linguistics.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. [Enriching Word Vectors with Subword Information](#). *Transactions of the Association for Computational Linguistics*, 5:135–146.
- Gemma Boleda. 2020. [Distributional semantics and linguistic theory](#). *Annual Review of Linguistics*, 6:213–234.
- Paula Czarnecka, Guy Emerson, and Ann Copestake. 2019. [Words are vectors, dependencies are matrices: Learning word embeddings from dependency graphs](#). In *Proceedings of the 13th International Conference on Computational Semantics - Long Papers*, pages 91–102, Gothenburg, Sweden. Association for Computational Linguistics.
- Guy Emerson. 2020. [What are the goals of distributional semantics?](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7436–7453, Online. Association for Computational Linguistics.
- Dan Flickinger, Stephan Oepen, and Gisle Ytrestøl. 2010. [WikiWoods: Syntacto-semantic annotation for English Wikipedia](#). In *Proceedings of the 7th International Conference on Language Resources and Evaluation (LREC)*, pages 1665–1671. European Language Resources Association (ELRA).
- Aurélie Herbelot and Marco Baroni. 2017. [High-risk learning: acquiring new word vectors from tiny data](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 304–309, Copenhagen, Denmark. Association for Computational Linguistics.
- Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Adriane Boyd. 2020. [spaCy: Industrial-strength natural language processing in Python](#).
- Mikhail Khodak, Nikunj Saunshi, Yingyu Liang, Tengyu Ma, Brandon Stewart, and Sanjeev Arora. 2018. [A la carte embedding: Cheap but effective induction of semantic feature vectors](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12–22, Melbourne, Australia. Association for Computational Linguistics.
- Angeliki Lazaridou, Marco Marelli, and Marco Baroni. 2017. [Multimodal word meaning induction from minimal exposure to natural text](#). *Cognitive Science*, 41 Suppl 4.
- Alessandro Lenci. 2018. [Distributional models of word meaning](#). *Annual review of Linguistics*, 4:151–171.
- Omer Levy and Yoav Goldberg. 2014. [Dependency-based word embeddings](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 302–308, Baltimore, Maryland. Association for Computational Linguistics.
- 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.
- Diego Marcheggiani and Ivan Titov. 2017. [Encoding sentences with graph convolutional networks for semantic role labeling](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1506–1515, Copenhagen, Denmark. Association for Computational Linguistics.

- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. [Distributed representations of words and phrases and their compositionality](#). In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26*, pages 3111–3119. Curran Associates, Inc.
- Ethan Perez, Douwe Kiela, and Kyunghyun Cho. 2021. [True few-shot learning with language models](#). In *Advances in Neural Information Processing Systems*.
- Magnus Sahlgren. 2008. [The distributional hypothesis](#). *Italian Journal of Disability Studies*, 20:33–53.
- Timo Schick and Hinrich Schütze. 2019. [Learning semantic representations for novel words: Leveraging both form and context](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6965–6973.
- Jeroen Van Hautte, Guy Emerson, and Marek Rei. 2019. [Bad form: Comparing context-based and form-based few-shot learning in distributional semantic models](#). In *Proceedings of the 2nd Workshop on Deep Learning Approaches for Low-Resource NLP (DeepLo 2019)*, pages 31–39, Hong Kong, China. Association for Computational Linguistics.
- Alexey Zobnin and Evgenia Elistratova. 2019. [Learning word embeddings without context vectors](#). In *Proceedings of the 4th Workshop on Representation Learning for NLP (RepLANLP-2019)*, pages 244–249, Florence, Italy. Association for Computational Linguistics.