

Conditional Generators of Words Definitions

Artyom Gadetsky

National Research University
Higher School of Economics
artygadetsky@yandex.ru

Ilya Yakubovskiy

Joom
yakubovskiy@joom.com

Dmitry Vetrov

National Research University
Higher School of Economics
Samsung-HSE Laboratory
vetrov@yandex.ru

Abstract

We explore recently introduced definition modeling technique that provided the tool for evaluation of different distributed vector representations of words through modeling dictionary definitions of words. In this work, we study the problem of word ambiguities in definition modeling and propose a possible solution by employing latent variable modeling and soft attention mechanisms. Our quantitative and qualitative evaluation and analysis of the model shows that taking into account words ambiguity and polysemy leads to performance improvement.

1 Introduction

Continuous representations of words are used in many natural language processing (NLP) applications. Using pre-trained high-quality word embeddings are most effective if not millions of training examples are available, which is true for most tasks in NLP (Kumar et al., 2016; Karpathy and Fei-Fei, 2015). Recently, several unsupervised methods were introduced to learn word vectors from large corpora of texts (Mikolov et al., 2013; Pennington et al., 2014; Joulin et al., 2016). Learned vector representations have been shown to have useful and interesting properties. For example, Mikolov et al. (2013) showed that vector operations such as subtraction or addition reflect semantic relations between words. Despite all these properties it is hard to precisely evaluate embeddings because analogy relation or word similarity tasks measure learned information indirectly.

Quite recently Noraset et al. (2017) introduced a more direct way for word embeddings evaluation. Authors suggested considering definition

modeling as the evaluation task. In definition modeling vector representations of words are used for conditional generation of corresponding word definitions. The primary motivation is that high-quality word embedding should contain all useful information to reconstruct the definition. The important drawback of Noraset et al. (2017) definition models is that they cannot take into account words with several different meanings. These problems are related to word disambiguation task, which is a common problem in natural language processing. Such examples of polysemantic words as “bank“ or “spring“ whose meanings can only be disambiguated using their contexts. In such cases, proposed models tend to generate definitions based on the most frequent meaning of the corresponding word. Therefore, building models that incorporate word sense disambiguation is an important research direction in natural language processing.

In this work, we study the problem of word ambiguity in definition modeling task. We propose several models which can be possible solutions to it. One of them is based on recently proposed Adaptive Skip Gram model (Bartunov et al., 2016), the generalized version of the original SkipGram Word2Vec, which can differ word meanings using word context. The second one is the attention-based model that uses the context of a word being defined to determine components of embedding referring to relevant word meaning. Our contributions are as follows: (1) we introduce two models based on recurrent neural network (RNN) language models, (2) we collect new dataset of definitions, which is larger in number of unique words than proposed in Noraset et al. (2017) and also supplement it with examples of the word usage (3) finally, in the experiment section we show that our models outperform previously proposed models and have the ability to generate

definitions depending on the meaning of words.

2 Related Work

2.1 Constructing Embeddings Using Dictionary Definitions

Several works utilize word definitions to learn embeddings. For example, Hill et al. (2016) use definitions to construct sentence embeddings. Authors propose to train recurrent neural network producing an embedding of the dictionary definition that is close to an embedding of the corresponding word. The model is evaluated with the reverse dictionary task. Bahdanau et al. (2017) suggest using definitions to compute embeddings for out-of-vocabulary words. In comparison to Hill et al. (2016) work, dictionary reader network is trained end-to-end for a specific task.

2.2 Definition Modeling

Definition modeling was introduced in Noraset et al. (2017) work. The goal of the definition model $p(D|w^*)$ is to predict the probability of words in the definition $D = \{w_1, \dots, w_T\}$ given the word being defined w^* . The joint probability is decomposed into separate conditional probabilities, each of which is modeled using the recurrent neural network with soft-max activation, applied to its logits.

$$p(D|w^*) = \prod_{t=1}^T p(w_t|w_{i < t}, w^*) \quad (1)$$

Authors of definition modeling consider following conditional models and their combinations: *Seed (S)* - providing word being defined at the first step of the RNN, *Input (I)* - concatenation of embedding for word being defined with embedding of word on corresponding time step of the RNN, *Gated (G)*, which is the modification of GRU cell. Authors use a character-level convolutional neural network (CNN) to provide character-level information about the word being defined, this feature vector is denoted as *(CH)*. One more type of conditioning referred to as *(HE)*, is hypernym relations between words, extracted using Hearst-like patterns.

3 Word Embeddings

Many natural language processing applications treat words as atomic units and represent them as continuous vectors for further use in machine

learning models. Therefore, learning high-quality vector representations is the important task.

3.1 Skip-gram

One of the most popular and frequently used vector representations is Skip-gram model. The original Skip-gram model consists of grouped word prediction tasks. Each task is formulated as a prediction of the word v given word w using their input and output representations:

$$p(v|w, \theta) = \frac{\exp(\text{in}_w^T \text{out}_v)}{\sum_{v'=1}^V \exp(\text{in}_w^T \text{out}_{v'})} \quad (2)$$

where θ and V stand for the set of input and output word representations, and dictionary size respectively. These individual prediction tasks are grouped in a way to independently predict all adjacent (with some sliding window) words $y = \{y_1, \dots, y_C\}$ given the central word x :

$$p(y|x, \theta) = \prod_j p(y_j|x, \theta) \quad (3)$$

The joint probability of the model is written as follows:

$$p(Y|X, \theta) = \prod_{i=1}^N p(y_i|x_i, \theta) \quad (4)$$

where $(X, Y) = \{x_i, y_i\}_{i=1}^N$ are training pairs of words and corresponding contexts and θ stands for trainable parameters.

Also, optimization of the original Skip-gram objective can be changed to a negative sampling procedure as described in the original paper or hierarchical soft-max prediction model (Mnih and Hinton, 2009) can be used instead of (2) to deal with computational costs of the denominator. After training, the input representations are treated as word vectors.

3.2 Adaptive Skip-gram

Skip-gram model maintains only one vector representation per word that leads to mixing of meanings for polysemantic words. Bartunov et al. (2016) propose a solution to the described problem using latent variable modeling. They extend Skip-gram to Adaptive Skip-gram (AdaGram) in a way to automatically learn the required number of vector representations for each word using Bayesian nonparametric approach. In comparison

with Skip-gram AdaGram assumes several meanings for each word and therefore keeps several vectors representations for each word. They introduce latent variable z that encodes the index of meaning and extend (2) to $p(v|z, w, \theta)$. They use hierarchical soft-max approach rather than negative sampling to overcome computing denominator.

$$p(v|z = k, w, \theta) = \prod_{n \in \text{path}(v)} \sigma(\text{ch}(n) \text{in}_{wk}^T \text{out}_n) \quad (5)$$

Here in_{wk} stands for input representation of word w with meaning index k and output representations are associated with nodes in a binary tree, where leaves are all possible words in model vocabulary with unique paths from the root to the corresponding leaf. $\text{ch}(n)$ is a function which returns 1 or -1 to each node in the $\text{path}(\cdot)$ depending on whether n is a left or a right child of the previous node in the path. Huffman tree is often used for computational efficiency.

To automatically determine the number of meanings for each word authors use the constructive definition of Dirichlet process via stick-breaking representation ($p(z = k|w, \beta)$), which is commonly used prior distribution on discrete latent variables when the number of possible values is unknown (e.g. infinite mixtures).

$$p(z = k|w, \beta) = \beta_{wk} \prod_{r=1}^{k-1} (1 - \beta_{wr}) \quad (6)$$

$$p(\beta_{wk}|\alpha) = \text{Beta}(\beta_{wk}|1, \alpha)$$

This model assumes that an infinite number of meanings for each word may exist. Providing that we have a finite amount of data, it can be shown that only several meanings for each word will have non-zero prior probabilities.

Finally, the joint probability of all variables in AdaGram model has the following form:

$$p(Y, Z, \beta|X, \alpha, \theta) = \prod_{w=1}^V \prod_{k=1}^{\infty} p(\beta_{wk}|\alpha) \cdot \prod_{i=1}^N [p(z_i|x_i, \beta) \prod_{j=1}^C p(y_{ij}|z_i, x_i, \theta)] \quad (7)$$

Model is trained by optimizing Evidence Lower Bound using stochastic variational inference (Hoffman et al., 2013) with fully factorized variational approximation of the posterior distribution $p(Z, \beta|X, Y, \alpha, \theta) \approx q(Z)q(\beta)$.

One important property of the model is an ability to disambiguate words using context. More formally, after training on data $D = \{x_i, y_i\}_{i=1}^N$ we may compute the posterior probability of word meaning given context and take the word vector with the highest probability.:

$$p(z = k|x, y, \theta) \propto p(y|x, z = k, \theta) \int p(z = k|\beta, x)q(\beta)d\beta \quad (8)$$

This knowledge about word meaning will be further utilized in one of our models as $\text{disambiguation}(x|y)$.

4 Models

In this section, we describe our extension to original definition model. The goal of the extended definition model is to predict the probability of a definition $D = \{w_1, \dots, w_T\}$ given a word being defined w^* and its context $C = \{c_1, \dots, c_m\}$ (e.g. example of use of this word). As it was motivated earlier, the context will provide proper information about word meaning. The joint probability is also decomposed in the conditional probabilities, each of which is provided with the information about context:

$$p(D|w^*, C) = \prod_{t=1}^T p(w_t|w_{i < t}, w^*, C) \quad (9)$$

4.1 AdaGram based

Our first model is based on original *Input (I)* conditioned on Adaptive Skip-gram vector representations. To determine which word embedding to provide as *Input (I)* we disambiguate word being defined using its context words C . More formally our *Input (I)* conditioning is turning in:

$$h_t = g([v^*; v_t], h_{t-1})$$

$$v^* = \text{disambiguation}(w^*|C) \quad (10)$$

where g is the recurrent cell, $[a; b]$ denotes vector concatenation, v^* and v_t are embedding of word being defined w and embedding of word w_t respectively. We refer to this model as *Input Adaptive (I-Adaptive)*.

4.2 Attention based

Adaptive Skip-gram model is very sensitive to the choice of concentration parameter in Dirichlet process. The improper setting will cause many

similar vectors representations with smoothed meanings due to theoretical guarantees on a number of learned components. To overcome this problem and to get rid of careful tuning of this hyper-parameter we introduce following model:

$$\begin{aligned} h_t &= g([a^*; v_t], h_{t-1}) \\ a^* &= v^* \odot \text{mask} \\ \text{mask} &= \sigma\left(W \frac{\sum_{i=1}^m \text{ANN}(c_i)}{m} + b\right) \end{aligned} \quad (11)$$

where \odot is an element-wise product, σ is a logistic sigmoid function and ANN is attention neural network, which is a feed-forward neural network. We motivate these updates by the fact, that after learning Skip-gram model on a large corpus, vector representation for each word will absorb information about every meaning of the word. Using soft binary mask dependent on word context we extract components of word embedding relevant to corresponding meaning. We refer to this model as *Input Attention (I-Attention)*.

4.3 Attention SkipGram

For attention-based model, we use different embeddings for context words. Because of that, we pre-train attention block containing embeddings, attention neural network and linear layer weights by optimizing a negative sampling loss function in the same manner as the original Skip-gram model:

$$\begin{aligned} &\log \sigma(v'_{w_O}{}^T v_{w_I}) \\ &+ \sum_{i=1}^k \mathbb{E}_{w_i \sim P_n(w)} [\log \sigma(-v'_{w_i}{}^T v_{w_I})] \end{aligned} \quad (12)$$

where v'_{w_O} , v_{w_I} and v'_{w_i} are vector representation of "positive" example, anchor word and negative example respectively. Vector v_{w_I} is computed using embedding of w_I and attention mechanism proposed in previous section.

5 Experiments

5.1 Data

We collected new dataset of definitions using [OxfordDictionaries.com \(2018\)](#) API. Each entry is a triplet, containing the word, its definition and example of the use of this word in the given meaning. It is important to note that in our data set words can have one or more meanings, depending on the corresponding entry in the Oxford Dictionary. Table 1 shows basic statistics of the new dataset.

Split	train	val	test
#Words	33,128	8,867	8,850
#Entries	97,855	12,232	12,232
#Tokens	1,078,828	134,486	133,987
Avg length	11.03	10.99	10.95

Table 1: Statistics of new dataset

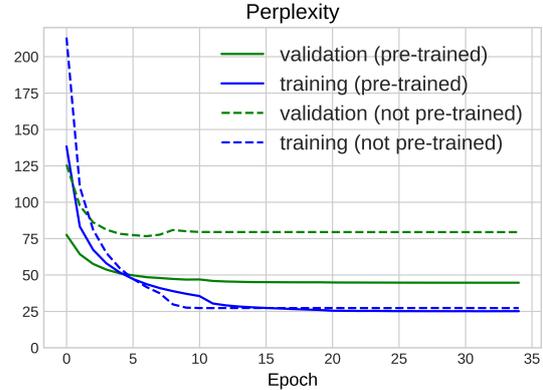


Figure 1: Perplexities of *S+I Attention* model for the case of pre-training (solid lines) and for the case when the model is trained from scratch (dashed lines).

5.2 Pre-training

It is well-known that good language model can often improve metrics such as BLEU for a particular NLP task (Jozefowicz et al., 2016). According to this, we decided to pre-train our models. For this purpose, WikiText-103 dataset (Merity et al., 2016) was chosen. During pre-training we set v^* (eq. 10) to zero vector to make our models purely unconditional. Embeddings for these language models were initialized by Google Word2Vec vectors¹ and were fine-tuned. Figure 1 shows that this procedure helps to decrease perplexity and prevents over-fitting. Attention Skip-gram vectors were also trained on the WikiText-103.

5.3 Results

Both our models are LSTM networks (Hochreiter and Schmidhuber, 1997) with an embedding layer. The attention-based model has own embedding layer, mapping context words to vector representations. Firstly, we pre-train our models using the procedure, described above. Then, we train them on the collected dataset maximizing log-likelihood objective using Adam (Kingma and Ba, 2014). Also, we anneal learning rate by

¹<https://code.google.com/archive/p/word2vec/>

Word	Context	Definition
star	she got star treatment	a person who is very important
star	bright star in the sky	a small circle of a celestial object or planet that is seen in a circle
sentence	sentence in prison	an act of restraining someone or something
sentence	write up the sentence	a piece of text written to be printed
head	the head of a man	the upper part of a human body
head	he will be the head of the office	the chief part of an organization, institution, etc
reprint	they never reprinted the famous treatise	a written or printed version of a book or other publication
rape	the woman was raped on her way home at night	the act of killing
invisible	he pushed the string through an inconspicuous hole	not able to be seen
shake	my faith has been shaken	cause to be unable to think clearly
nickname	the nickname for the u.s. constitution is 'old ironsides '	a name for a person or thing that is not genuine

Table 2: Examples of definitions generated by $S + I$ -Attention model for the words and contexts from the test set.

Model	PPL	BLEU
$S+G+CH+HE$ (1)	45.62	11.62 ± 0.05
$S+G+CH+HE$ (2)	46.12	-
$S+G+CH+HE$ (3)	46.80	-
$S + I$ -Adaptive (2)	46.08	11.53 ± 0.03
$S + I$ -Adaptive (3)	46.93	-
$S + I$ -Attention (2)	43.54	12.08 ± 0.02
$S + I$ -Attention (3)	44.9	-

Table 3: Performance comparison between best model proposed by Noraset et al. (2017) and our models on the test set. Number in brackets means number of LSTM layers. BLEU is averaged across 3 trials.

a factor of 10 if validation loss doesn't decrease per epochs. We use original Adaptive Skip-gram vectors as inputs to $S+I$ -Adaptive, which were obtained from the official repository². We compare different models using perplexity and BLEU score on the test set. BLEU score is computed only for models with the lowest perplexity and only on the test words that have multiple meanings. The results are presented in Table 3. We see that both models that utilize knowledge about meaning of the word have better performance than the competing one. We generated definitions using $S + I$ -Attention model with simple temperature sampling

algorithm ($\tau = 0.1$). Table 2 shows the examples. The source code and dataset will be freely available³.

6 Conclusion

In the paper, we proposed two definition models which can work with polysemantic words. We evaluate them using perplexity and measure the definition generation accuracy with BLEU score. Obtained results show that incorporating information about word senses leads to improved metrics. Moreover, generated definitions show that even implicit word context can help to differ word meanings. In future work, we plan to explore individual components of word embedding and the mask produced by our attention-based model to get a deeper understanding of vectors representations of words.

Acknowledgments

This work was partly supported by Samsung Research, Samsung Electronics, Sberbank AI Lab and the Russian Science Foundation grant 17-71-20072.

²<https://github.com/sbos/AdaGram.jl>

³<https://github.com/agadetsky/pytorch-definitions>

References

- Dzmitry Bahdanau, Tom Bosc, Stanislaw Jastrzebski, Edward Grefenstette, Pascal Vincent, and Yoshua Bengio. 2017. Learning to compute word embeddings on the fly. *arXiv preprint arXiv:1706.00286*.
- Sergey Bartunov, Dmitry Kondrashkin, Anton Osokin, and Dmitry Vetrov. 2016. Breaking sticks and ambiguities with adaptive skip-gram. In *Artificial Intelligence and Statistics*, pages 130–138.
- Felix Hill, KyungHyun Cho, Anna Korhonen, and Yoshua Bengio. 2016. Learning to understand phrases by embedding the dictionary. *Transactions of the Association for Computational Linguistics*, 4:17–30.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Comput.*, 9(8):1735–1780.
- Matthew D. Hoffman, David M. Blei, Chong Wang, and John Paisley. 2013. [Stochastic variational inference](#). *Journal of Machine Learning Research*, 14:1303–1347.
- Armand Joulin, Edouard Grave, Piotr Bojanowski, Matthijs Douze, Herve Jégou, and Tomas Mikolov. 2016. Fasttext.zip: Compressing text classification models. *arXiv preprint arXiv:1612.03651*.
- Rafal Jozefowicz, Oriol Vinyals, Mike Schuster, Noam Shazeer, and Yonghui Wu. 2016. Exploring the limits of language modeling. *arXiv preprint arXiv:1602.02410*.
- Andrej Karpathy and Li Fei-Fei. 2015. Deep visual-semantic alignments for generating image descriptions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3128–3137.
- Diederik P. Kingma and Jimmy Ba. 2014. [Adam: A method for stochastic optimization](#). *CoRR*, abs/1412.6980.
- Ankit Kumar, Ozan Irsoy, Peter Ondruska, Mohit Iyyer, James Bradbury, Ishaan Gulrajani, Victor Zhong, Romain Paulus, and Richard Socher. 2016. Ask me anything: Dynamic memory networks for natural language processing. In *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 1378–1387. PMLR.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2016. Pointer sentinel mixture models. *arXiv preprint arXiv:1609.07843*.
- 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.
- Andriy Mnih and Geoffrey E Hinton. 2009. A scalable hierarchical distributed language model. In *Advances in Neural Information Processing Systems 21*, pages 1081–1088. Curran Associates, Inc.
- Thanapon Noraset, Chen Liang, Larry Birnbaum, and Doug Downey. 2017. Definition modeling: Learning to define word embeddings in natural language. In *31st AAAI Conference on Artificial Intelligence, AAAI 2017*, pages 3259–3266. AAAI press.
- OxfordDictionaries.com. 2018. Oxford University Press.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.