



WSD

Loïc Vial, Benjamin Lecouteux, Didier Schwab

► To cite this version:

| Loïc Vial, Benjamin Lecouteux, Didier Schwab. WSD. [Research Report] LIG. 2018. hal-01940293

HAL Id: hal-01940293

<https://hal.science/hal-01940293>

Submitted on 30 Nov 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Improving the Coverage and the Generalization Ability of Neural Word Sense Disambiguation through Hypernymy and Hyponymy Relationships

Loïc Vial Benjamin Lecouteux Didier Schwab

Abstract

In Word Sense Disambiguation (WSD), the predominant approach generally involves a supervised system trained on sense annotated corpora. The limited quantity of such corpora however restricts the coverage and the performance of these systems. In this article, we propose a new method that solves these issues by taking advantage of the knowledge present in WordNet, and especially the hypernymy and hyponymy relationships between synsets, in order to reduce the number of different sense tags that are necessary to disambiguate all words of the lexical database. Our method leads to state of the art results on most WSD evaluation tasks, while improving the coverage of supervised systems, reducing the training time and the size of the models, without additional training data. In addition, we exhibit results that significantly outperform the state of the art when our method is combined with an ensemble technique and the addition of the WordNet Gloss Tagged as training corpus.

1 Introduction

Word Sense Disambiguation (WSD) is a task which aims to clarify a text by assigning to each of its words the most suitable sense labels, given a predefined sense inventory.

Various approaches have been proposed to achieve WSD, and they are generally ordered by the type and the quantity of resources they use:

- **Knowledge-based** methods rely on dictionaries, lexical databases, thesauri or knowledge graphs as main resources, and use algorithms such as lexical similarity measures (Lesk, 1986) or graph-based measures (Moro et al., 2014).

- **Supervised** methods, on the other hand, exploit sense annotated corpora as training instances that can be used by a multi-class classifier such as SVM (Chan et al., 2007; Zhong and Ng, 2010), or more recently by a neural network (Kågebäck and Salomonsson, 2016).
- **Semi-supervised** methods generally use unannotated data to artificially increase the quantity of sense annotated data and hence improve supervised methods (Yuan et al., 2016).

Supervised methods are by far the most predominant as they generally offer the best results in evaluation campaigns (for instance (Navigli et al., 2007)). State of the art classifiers used to combine a set of specific features such as the parts of speech tags of surrounding words, local collocations (Zhong and Ng, 2010) and pretrained word embeddings (Iacobacci et al., 2016), but they are now replaced by recurrent neural networks which learn their own representation of words (Raganato et al., 2017b; Le et al., 2018).

One of the major bottleneck of supervised systems is the restricted quantity of manually sense annotated corpora. Indeed, while the lexical database WordNet (Miller, 1995), the sense inventory of reference used in most works on WSD, contains more than 200 000 different word-sense pairs¹, the SemCor (Miller et al., 1993), the corpus which is used the most in the training of supervised systems, only represents approximately 34 000 of them.

Many works try to leverage this problem by creating new sense annotated corpora, either automatically (Pasini and Navigli, 2017),

¹<https://wordnet.princeton.edu/documentation/wnstats7wn>

semi-automatically (Taghipour and Ng, 2015), or through crowdsourcing (Yuan et al., 2016), but in this work, the idea is to solve this issue by taking advantage of one of the multiple semantic relationships between senses included in WordNet: the hypernymy and hyponymy relationships. Our method is based on three observations:

1. A sense, its hypernym and its hyponyms share a common idea or concept, but on different levels of abstraction.
2. In general, a word can be disambiguated using the hypernyms of its senses, and not necessarily the senses themselves.
3. Consequently, we do not need to know every sense of WordNet to disambiguate all words of WordNet.

Contributions: We propose a method for reducing the vocabulary of senses of WordNet by selecting the minimal set of senses required for differentiating the meaning of every word. By using this technique, and converting the sense tags present in sense annotated corpora to the most generalized sense possible, we are able to greatly improve the coverage and the generalization ability of supervised systems.

We start by presenting the state of the art of supervised neural architectures for word sense disambiguation, then we describe our new method for sense vocabulary reduction. Our method is then evaluated by measuring its contribution to a state of the art neural WSD system evaluated on classic WSD evaluation campaigns.

The code for using our system or reproducing our results is available at the following URL:

<https://github.com/getalp/disambiguate>

2 Neural Word Sense Disambiguation

The neural approaches for WSD fall into two categories: approaches based on a neural model that learns to classify a sense directly, and approaches based on a neural language model that learns to predict a word, and is then used to find the closest sense to the predicted word.

2.1 Language Model Based WSD

The core of these approaches is a powerful neural language model able to predict a word with consideration for the words surrounding it, thanks to a recurrent neural network trained on a massive quantity of unannotated data. The main works that implement these kind of model are Yuan et al. (2016) and Le et al. (2018).

Once the language model is trained, its predictions are used to produce sense vectors as the average of the word vectors predicted by the language model in the places where the words are sense annotated.

At test time, the language model is used to predict a vector according to the surrounding context, and the sense closest to the predicted vector is assigned to each word.

These systems have the advantage of bypassing the problem of the lack of sense annotated data by concentrating the power of abstraction offered by recurrent neural networks on a good quality language model trained in an unsupervised manner. However, sense annotated corpora are still indispensable to construct the sense vectors, and the quantity of data needed for training the language model (100 billion tokens for Yuan et al. (2016), 2 billion tokens for Le et al. (2018)) makes these systems more difficult to train than those relying on sense annotated data only.

2.2 Classification Based WSD

In these systems, the main neural network directly classifies and attributes a sense to each input word. Sense annotations are simply seen as tags put on every word, like a POS-tagging task for instance.

These models are more similar to classical supervised models such as Chan et al. (2007), except that the input features are not manually selected, but trained as part of the neural network (using pre-trained word embeddings or not).

In addition, we can distinguish two separate branches of these types of neural networks:

1. Those in which we have several distinct and small neural networks (or classifiers) for every different word in the dictionary (Iacobacci et al., 2016; Kågebäck and Salomonsson, 2016), each of them being able

to manage a particular word and its particular senses. For instance, one of the classifiers is specialized into choosing between the four possible senses of the noun “mouse”. This type of approaches is particularly fitted for the lexical sample tasks, where a small and finite set of very ambiguous words have to be sense annotated in several contexts, but it can also be used in all-words word sense disambiguation tasks.

2. Those in which we have a bigger and unique neural network that is able to manage all different words and assign a sense in the set of all existing sense in the dictionary used (Raganato et al., 2017b).

The advantage of the first branch of approaches is that in order to disambiguate a word, limiting our choice to one of its possible senses is computationally much easier than searching through all the senses of all words. To put things in perspective, the number of senses of each word in WordNet ranges from 1 to 59, whereas the total number of senses considering all words is 206 941.

The other approach however has an interesting property: all senses reside in the same vector space and hence share features in the hidden layers of the network. This allows the model to predict a common sense for two different words (i.e. synonyms), but it also offers the possibility to predict a sense for a word not present in the dictionary (e.g. neologism, spelling mistake...), and let the user or the underlying system to decide afterwards what to do with this prediction.

In practice, this ability of merging multiple sense tags together is especially useful when working with WordNet: indeed, this lexical database is based on the notion of synonym sets or “synsets”, group of senses with the same meaning and definition. Disambiguating with synset tags instead of sense tags is a common practice (Yuan et al., 2016; Le et al., 2018), as it effectively decreases the output vocabulary of the classifier that considers all senses in WordNet from 206 941 to 117 659, and one can unambiguously retrieve the sense tag given a synset tag and the tagged word (because every sense of a word belong to a different synset).

In this work, we go further into this direction and we present a method based on the hy-

ponymy and hyponymy relationships present in WordNet, in order to merge synset tags together and reduce even more the output vocabulary of such neural WSD systems.

3 Sense Vocabulary Reduction

We can draw three issues of the current situation regarding supervised WSD systems:

1. The training of systems that directly predict a tag in the set of all WordNet senses becomes slower and take more memory the larger the output vocabulary is. This output layer size going up to 206 941 if we consider all word-senses, and 117 659 if we consider all synsets.
2. Due to the small number of manually sense annotated corpora available, a target word may never be observed during the training, and therefore the system would not be able to annotate it.
3. For the same reason, a word may have been observed, but not all of its senses. In this case the system is able to annotate the word, but if the expected sense has never been observed, the output will be wrong, regardless of the architecture of the supervised system.

In the SemCor (Miller et al., 1993) for instance, the largest manually sense annotated corpus available, words are annotated with 33 760 different sense keys, which corresponds to approximately 16% of the sense inventory of WordNet².

Grouping together multiple senses is hence a good way to overcome all these issues: by considering that multiple tags refer in fact to the same concept, the output vocabulary decreases, the ability of the trained system to generalize improves, and also its coverage.

Moreover, it reflects more accurately our intuition of what a sense is: clearly the notions of “tree” (with a trunk and leaves, not the mathematical graph) and “plant” (the living organism, not the industrial building) forms a group in our mind such that observing one sense in a context should help disambiguating the other.

²<https://wordnet.princeton.edu/documentation/wnstats7wn>

3.1 From Senses to Synsets: A First Sense Vocabulary Reduction

WordNet is a lexical database organized in sets of synonyms called synsets. A synset is technically a group of one or more word-senses that have the same definition and consequently the same meaning. For instance, the first senses of “eye”, “optic” and “oculus” all refer to a common synset which definition is “the organ of sight”.

Training a WSD supervised system to predict synset tags instead of word-sense tags is a common practice (Yuan et al., 2016; Le et al., 2018), and it can be seen as a form of output vocabulary reduction based on the knowledge that is present in WordNet.

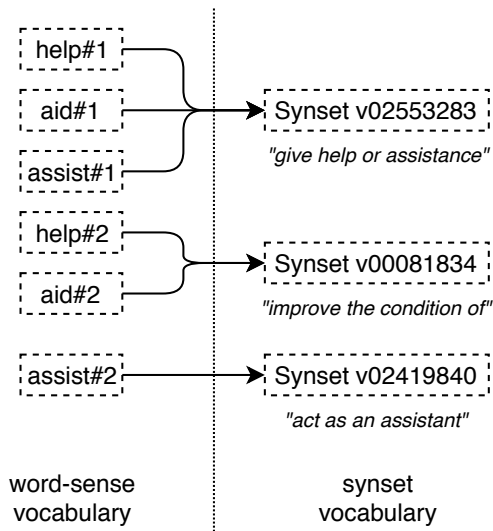


Figure 1: Word-sense to synset vocabulary reduction applied on the first three senses of the words “help”, “aid” and “assist”

Illustrated in Figure 1, the word-sense to synset vocabulary reduction clearly helps to improve the coverage of supervised systems. Indeed, if the verb “help” is observed in the annotated data in its first sense, and consequently with the tag “v02553283”, the context surrounding the target word can be used to later annotate the verb “assist” or “aid” with the same valid tag.

Once applied, the number of different labels needed to cover all senses of WordNet drops from 206 941 to 117 659 (approximately 43% of reduction), and considering the SemCor, the corpus contains 26 215 different synsets, which accounts now for 22% of this total. The vocab-

ulary size was reduced and the coverage improved.

Going a little further, other information from WordNet can help the system to generalize. In the next section, we describe a new method for taking advantage of the hypernymy and hyponymy relationships in order to accomplish this same idea.

3.2 Sense Vocabulary Reduction through Hypernymy and Hyponymy Relationships

According to Polguère (2003), hypernymy and hyponymy are two semantic relationships which correspond to a particular case of sense inclusion: the hyponym of a term is a specialization of this term, whereas its hypernym is a generalization. For instance, a “mouse” is a type of “rodent” which is in turn a type of “animal”.

In WordNet, these relationships bind nearly every nouns³ together in a tree structure that goes from the generic root, the node “entity” to the most specific leaves, for instance the node “white-footed mouse”. These relationships are also present on several verbs: so for instance “add” is a way of “compute” which in turn is a way of “reason” which is a way of “think”.

For the sake of WSD, just like grouping together the senses of a same synset helps to better generalize, we hypothesize that grouping together the synsets of a same hypernymy relationship also helps in the same way.

The general idea of our method is that the most specialized concepts in WordNet are often superfluous in order to perform WSD.

Indeed, consider a small subset of WordNet that only consists of the word “mouse”, its first sense (the small rodent), its fourth sense (the electronic device), and all of their hypernyms. This is illustrated in Figure 2. We can see that every concept that is more specialized than the concepts “artifact” and “living_thing” could be removed, and we could map every tag of “mouse#1” to the tag of “living_thing#1” and we could still be able to disambiguate this word, but with a benefit: all other “living things” and animals in the sense annotated data could be tagged with the same

³Only 2% of the polysemous nouns of WordNet are not part of this hierarchy.

sense, give examples of what is an animal and then show how to differentiate the small rodent to the hand-operated electronic device.

In order to achieve this goal of mapping every sense to its most generic sense still allowing to differentiate the meanings of the words, we have to consider certain difficulties that are not present with the word-sense to synset vocabulary reduction.

First, contrary to the synonymy relationship which is symmetric (i.e. if A is a synonym of B then B is a synonym of A), the hypernymy relationship is not. For instance, all mice are animals, but not all animals are mice.

In addition, two different senses of a word necessarily have two different synsets, but they may have the same direct hypernym, and they generally have the same inherited hypernym at a certain point. For instance, we can distinguish the sense 1 of “mouse” which is a type of “animal” from the sense 4 which is a type of “electronic device”, but we cannot distinguish them if we go too far into the hypernymy hierarchy, because both of them are a type of “physical entity”.

Finally, we could think of removing a synset from the vocabulary of WordNet because it is not useful locally (from the point of view of a specific word), but it could be necessary to differentiate the meanings of another word.

Our method thus works in two steps:

1. We mark as “necessary” all synsets that are the lowest nodes of the hypernymy hierarchies of the senses of all word that can still allow to discriminate the different senses of the word.
2. We transform our sense vocabulary by mapping every synset to the lowest synset in its hypernymy hierarchy that is marked as “necessary”.

The result of this method is that the most specific synsets of the tree that are not useful for discriminating are automatically removed from the vocabulary. In other words, the set of synsets that is left in the vocabulary is the smallest subset of all synsets that are necessary to distinguish every sense of every word of WordNet.

When applied on WordNet, the number of synset in the vocabulary now drops from

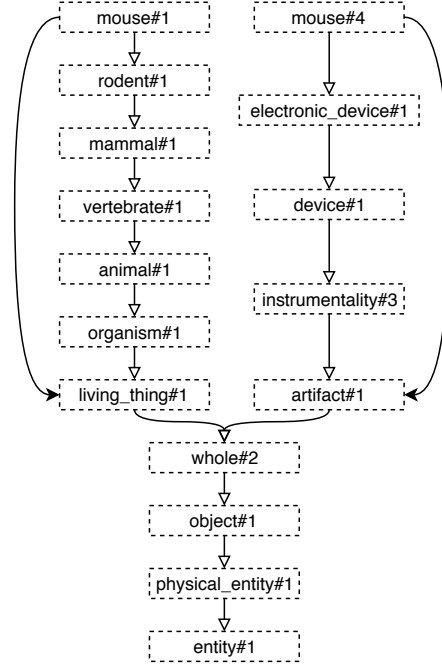


Figure 2: Sense reduction through hypernymy hierarchy applied on the first and fourth sense of the word “mouse” (some nodes are missing for clarity)

117659 to 39147 (approximately 66% of reduction), and applied on the SemCor, it now contains 12779 different synsets, which counts for 32% of coverage. Again, the vocabulary size has drastically decreased, and the coverage really improved. Note that if we narrow down our computation to consider only polysemic words in WordNet, the full vocabulary of all reduced synsets of WordNet is 23148, and the SemCor contains 9461 of them represented, and that is a coverage of approximately 40%.

4 Experiments

In order to evaluate our vocabulary reduction method, we applied it on a classification based neural network (subsection 2.2) capable of classifying a word in all possible synsets of WordNet. Our architecture is very similar to Raganato et al. (2017b)’s BiLSTM model except for the input and output vocabulary used. Indeed, in their system, they have chosen to include their input vocabulary in their output vocabulary, so their network is able to predict both a sense tag when the target word has an entry in WordNet (nouns, verbs, adjectives and averbs), and a word tag for every other word

(pronouns, articles, etc.). In our architecture, we chose to only predict sense tags, in order to keep the output vocabulary the smallest possible.

Then we systematically trained two models:

1. A baseline model that predicts a tag belonging to all the synset tags seen during training (thus using the common word-sense to synset vocabulary reduction).
2. A second system trained under the same conditions, but with our vocabulary reduction through hypernyms algorithm applied on the training corpus.

4.1 Neural Architecture

The architecture of our neural network relies on 3 layers :

1. The input layer, which takes directly the words in a vector form, from a pre-trained word embeddings model.
2. The hidden layer, composed of bidirectional LSTM units (Hochreiter and Schmidhuber, 1997).
3. The output layer, which represents for each word in input, a probability distribution over all senses in the output vocabulary used, thanks to a classical softmax function.

The cost function to minimize during the training is the cross entropy between the output layer and a one-hot vector, i.e. a vector for which all coordinates are set to 0 except for the coordinate at the index of the target sense which is 1. In consequence, the cost function is $-\log q[s]$, where $q[s]$ is the output of the network at the index s of the target sense.

Our model always predicts a sense in output, for every input word, even for words that do not convey directly a meaning (e.g. stop-words, articles, etc.) or words that were not annotated in the training set. However, we assign a special tag *<skip>* to these cases, allowing us to ignore the predictions made by the model and to not take it into account during the back-propagation step of the training.

This behavior is the main difference between our architecture and the one introduced by (Raganato et al., 2017b). In their model, the

gradient is computed over all words of a sentence, and those that do not have a sense in WordNet are annotated with their surface form.

In input of our network, we used the GloVe vectors (Pennington et al., 2014) pretrained on Wikipedia 2014 and Gigaword 5⁴. The dimension of the vectors is 300, the vocabulary size is 400 000 and all words are lowercased. These vectors are also used as input in the network described by (Kågeback and Salomonsson, 2016).

For the hidden layer of recurrent units, we chose LSTM cells of size 1000 for each direction (so 2000 in total). This is approximately the same size that was used in (Raganato et al., 2017b) (1024 per direction) and (Yuan et al., 2016) (a single layer of size 2048).

Finally, we applied the regularization method Dropout (Srivastava et al., 2014) between the hidden layer and the output layer, with a parameter set to 50%, in order to avoid overfitting during the training and to make the model more robust.

We implemented our neural network using PyTorch⁵, and our code is available at the following URL:

<https://github.com/getalp/disambiguate>

4.2 Training

We compared our sense vocabulary reduction method on two training sets: The first is the SemCor (Miller et al., 1993), the most popular corpus that is used for training most WSD supervised systems. The second is the concatenation of the SemCor and the WordNet Gloss Tagged⁶. The latter is a corpus distributed as part of WordNet since its version 3.0, and it consists of all the definitions (glosses) of every synset of WordNet, with every word manually or semi-automatically sense annotated. We used the version of these corpora given as part of the UFSAC 2.1 resource⁷ (Vial et al., 2018), a set of gathered publicly available sense annotated corpora converted into a clean and unified format.

We performed every training for 20 epochs.

⁴<https://nlp.stanford.edu/projects/glove/>

⁵<https://pytorch.org>

⁶<http://wordnetcode.princeton.edu/glosstag-files/glosstag.shtml>

⁷<https://github.com/getalp/UFSAC>

System	SE2	SE3	SE07 (17)	SE13 (12)	SE15 (13)	ALL	SE07 (07)
SemCor, baseline	91.15	96.76	97.58	91.06	94.78	93.23	92.84
SemCor, vocabulary reduced	98.03	99.19	99.78	99.15	98.39	98.75	98.85
SemCor+WNGT, baseline	97.81	98.92	99.34	97.63	99.34	98.26	98.45
SemCor+WNGT, vocabulary reduced	99.74	99.95	100	99.76	99.91	99.83	99.91

Table 1: Coverage of supervised systems based on the training corpus and if the vocabulary reduction algorithm is applied or not. Numbers are the percentage of words that are observed during training and hence can be annotated.

That is, the whole training set has been read 20 times. At the beginning of each epoch we shuffled the training set. We evaluated our model at the end of every epoch on a development set, and we kept only the one which obtained the best F1 WSD score. The development set was composed of 4 000 random sentences taken from the WordNet Gloss Tagged for the models trained on the SemCor, and 4 000 random sentences extracted from the training set for the other models.

We trained with mini-batches of 100 sentences, truncated to 80 words, and padded with zero vectors from the end, and we used Adam (Kingma and Ba, 2014), with the same default parameters described in their article as the optimization method, except for the learning rate that we set to 0.0001 (10 times smaller than the default value).

All models have been trained on Nvidia’s Titan X GPUs. The approximate training times of individual models, depending on the training corpus and if the vocabulary reduction method was applied, are displayed in the following table:

Model	SemCor	SemCor+WNGT
baseline	50mn	360mn (6h)
vocabulary reduced	30mn	160mn (2h40)

4.3 Disambiguation

In order to disambiguate an input sequence of words using the trained model, we followed the following steps:

First, each word of the sequence is lower-cased and transformed into a vector using the pre-trained word embeddings model, then the sequence of vectors is given as input to our

model.

Then, we annotate each word with the one among its possible senses which has the maximum probability. We first map each sense to its synset in the case of the baseline model, or map each sense to its reduced synset, in the case of the sense vocabulary reduced model, according to the method described in subsection 3.2, then we select the one which has the maximum value in the output layer of the model.

Finally, if no sense is assigned, because no instance of the word has been observed in the training data, a back-off is performed. We chose the most common one which is to assign the first sense in WordNet.

4.4 Evaluation

We evaluated our models on all evaluation corpora commonly used in WSD, that is the WSD tasks of the evaluation campaigns SensEval/SemEval. We used the fine-grained evaluation corpora from the evaluation framework of Raganato et al. (2017a), which consists of SensEval 2 (Edmonds and Cotton, 2001), SensEval 3 (Snyder and Palmer, 2004), SemEval 2007 task 17 (Pradhan et al., 2007), SemEval 2013 task 12 (Navigli et al., 2013) and SemEval 2015 task 13 (Moro and Navigli, 2015), as well as their “ALL” corpus consisting of the concatenation of all previous ones. We also compared our result on the coarse-grained task 7 of SemEval 2007 (Navigli et al., 2007) which is not present in the evaluation framework. We used the version of these corpora from the UFSAC 2.1 resource⁸, the sense inventory used for the sense annotations is WordNet 3.0.

⁸<https://github.com/getalp/UFSAC>

System	SE2	SE3	SE07 (17)	SE13 (12)	SE15 (13)	ALL	SE07 (07)
Our system (SemCor, baseline)	73.50 (± 0.26)	70.86 (± 0.42)	62.47 (± 1.05)	67.55 (± 0.41)	71.60 (± 0.52)	70.51 (± 0.16)	83.77 (± 0.28)
Our system (SemCor, vocabulary reduced)	72.55 (± 0.50)	70.39 (± 0.41)	61.46 (± 1.28)	70.77 (± 0.55)	71.32 (± 0.49)	70.77 (± 0.21)	84.48 (± 0.27)
Our system (SemCor+WNGT, baseline)	74.40 (± 0.51)	70.82 (± 0.49)	62.48 (± 1.29)	70.82 (± 0.58)	74.46 (± 0.71)	71.93 (± 0.35)	84.95 (± 0.42)
Our system (SemCor+WNGT, vocabulary reduced)	74.28 (± 0.40)	69.58 (± 0.50)	64.46 (± 1.31)	71.47 (± 0.79)	73.71 (± 0.68)	71.79 (± 0.29)	85.79 (± 0.26)
(Yuan et al., 2016) (SemCor, LSTM)	73.6	69.2	64.2	67.0	72.1	-	82.8
(Yuan et al., 2016) (SemCor, LSTM + LP)	73.8	71.8	63.5	69.5	72.6	-	83.6
(Raganato et al., 2017b) (SemCor, BLSTM)	71.4	68.8	*61.8	65.6	69.2	68.9	-
(Raganato et al., 2017b) (SemCor, BLSTM + att. + LEX + POS)	72.0	69.1	* 64.8	66.9	71.5	69.9	83.1
(Iacobacci et al., 2016) (IMS+emb) (Results from Raganato et al. (2017b))	72.2	70.4	62.6	65.9	71.5	70.1	81.9
First sense baseline	65.6	66.0	54.5	63.8	67.1	65.5	78.9

Table 2: F1 scores (%) obtained by our systems against the state of the art on the English WSD tasks of the evaluation campaigns SensEval 2 (SE2), SensEval 3 (SE3), SemEval 2007 (SE07) task 7 and 17, SemEval 2013 (SE13) task 12, SemEval 2015 (SE15) task 13 and the corpus composed of the concatenation of all previous ones (ALL) except SE07 task 7. Results in **bold** are the best results from using the sense vocabulary reduction or not. Results in **red** are to our knowledge the best results obtained on the task. Our results are the mean scores of 20 individual systems, with the standard deviation given in parenthesis. Results prefixed by a star (*) was obtained on the development corpus used during the training.

System (ensemble)	SE2	SE3	SE07 (17)	SE13 (12)	SE15 (13)	ALL	SE07 (07)
Train on SemCor, baseline	73.93	70.81	63.74	67.27	72.27	70.77	83.77
Train on SemCor, vocabulary reduced	73.05	70.59	61.32	71.23	71.60	71.18	84.65
Train on SemCor+WNGT, baseline	74.72	71.08	62.86	71.11	74.83	72.23	85.01
Train on SemCor+WNGT, vocabulary reduced	75.15	70.11	66.81	72.63	74.46	72.74	86.02

Table 3: F1 scores (%) obtained by our system with an ensemble of 20 models trained separately. Results in **bold** are the best results from all our systems. Results in **red** are to our knowledge the best results obtained on the task.

For each evaluation, we trained 20 separated models, and we give two scores: First, the mean of the F1 scores obtained by the models, along with its standard deviation. Then, the score obtained by an ensemble of the models. For the ensemble, we averaged the predictions of all individual models through a geometric mean, a common practice that is used for instance in machine translation (Sutskever et al., 2014; Gehring et al., 2017).

4.5 Results

The scores obtained by our systems using a single trained model compared to the state-of-the-art systems (Yuan et al., 2016; Raganato et al., 2017b; Iacobacci et al., 2016), along with the first sense baseline are present in table 2. The scores obtained by our ensemble of models are given in Table 3.

In subsection 3.2, we showed that our vocabulary reduction method improves the coverage of supervised systems over all WordNet vocabulary. In Table 1, we can see that this coverage improvement holds true on the evaluation tasks, for both training sets. On the total of 7253 words to annotate for the corpus “ALL”, the baseline system trained on the SemCor only is incapable of annotating 491 of them, and with the vocabulary reduction applied this number drops to 91. When adding the WordNet Gloss Tagged to the training set, this number is 126 for the baseline system, and with the vocabulary reduction, only 12 words cannot be annotated.

Now if we look at the results in Table 2, the difference of scores obtained by our system using the sense vocabulary reduction or not is overall not significant (regarding the “ALL” column). However we can notice a very large gap on the SemEval 2013 task, especially when the SemCor is used alone for training. This can be explained by the fact that this corpus is only composed of nouns, and our method for vocabulary reduction targets this part of speech principally. This is also the task where the coverage was improved the most by our method, as it can be seen in Table 1.

In comparison with the other works, our systems trained on the SemCor alone expose results comparable with the best system of Yuan et al. (2016), which is trained on the same corpus and augmented with a semi-supervised

method. When we add the WordNet Gloss Tagged to the training data however, we obtain systematically state of the art results on all tasks except on SensEval 3. Once again, the sense reduction method does not consistently improves or decreases the score on every task, and in overall (task “ALL”), the result is roughly the same as without sense reduction applied.

Finally, in Table 3 we show the results of our system ensembling 20 models by averaging the output of their last layer. As we can see, ensembling is a very efficient method in WSD as it improves systematically all our results. Interestingly, with ensembles, the scores are significantly higher when applying the vocabulary reduction algorithm. One possible interpretation is that individual models might be more frequently “lost” in the sense that with the sense vocabulary reduction applied, a lot of words are annotated with the same tag, and it can make the trained model “unsure” about the decisions it make. Ensemble of models tends to prevent this problem by favoring the most probable decisions of the models.

5 Conclusion

In this paper, we presented a new method that improves the coverage and the capacity of generalization of most supervised systems, by narrowing down the number of different sense in WordNet in order to keep only the senses that are essential for differentiating the meaning of all words present in the lexical database.

By considering that a same sense tag may be applied to many different words, this method also captures in a intuitive way a better representation of what is a sense for the task of word sense disambiguation.

We implemented a state of the art neural network for WSD and we showed that this method really improves the overall results using the same training corpus, especially when making ensembles of models, and also especially when disambiguating nouns.

We trained two sets of systems, one relying on the SemCor alone, and one with the addition of the WordNet Gloss Tagged corpus, and in this last configuration we obtained our best results that significantly outperform the state of the art on most WSD tasks.

With the combination of our sense vocabulary reduction method through the hypernymy hierarchy of WordNet and the addition of the WordNet Gloss Tagged to the set of training corpora, the coverage of our supervised system is almost 100% on most WSD tasks, and so this provides a solid alternative to the automatic or semi-automatic creation of sense annotated corpora, and this nearly eliminates the need for a first sense backoff in WSD supervised systems.

References

- Yee Seng Chan, Hwee Tou Ng, and Zhi Zhong. 2007. [Nus-pt: Exploiting parallel texts for word sense disambiguation in the english all-words tasks](#). In *Proceedings of the 4th International Workshop on Semantic Evaluations*, SemEval '07, pages 253–256, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Philip Edmonds and Scott Cotton. 2001. [Senseval-2: Overview](#). In *The Proceedings of the Second International Workshop on Evaluating Word Sense Disambiguation Systems*, SENSEVAL '01, pages 1–5, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N. Dauphin. 2017. [Convolutional sequence to sequence learning](#). In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1243–1252, International Convention Centre, Sydney, Australia. PMLR.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. [Long short-term memory](#). *Neural Computation*, 9(8):1735–1780.
- Ignacio Iacobacci, Mohammad Taher Pilehvar, and Roberto Navigli. 2016. [Embeddings for word sense disambiguation: An evaluation study](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 897–907, Berlin, Germany. Association for Computational Linguistics.
- Mikael Kågebäck and Hans Salomonsson. 2016. Word sense disambiguation using a bidirectional lstm. In *5th Workshop on Cognitive Aspects of the Lexicon (CogALex)*. Association for Computational Linguistics.
- Diederik P. Kingma and Jimmy Ba. 2014. [Adam: A method for stochastic optimization](#). *CoRR*, abs/1412.6980.
- Minh Le, Marten Postma, Jacopo Urbani, and Piek Vossen. 2018. [A deep dive into word sense disambiguation with lstm](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 354–365. Association for Computational Linguistics.
- Michael Lesk. 1986. Automatic sense disambiguation using mrd: how to tell a pine cone from an ice cream cone. In *Proceedings of SIGDOC '86*, pages 24–26, New York, NY, USA. ACM.
- George A Miller. 1995. Wordnet: a lexical database for english. *Communications of the ACM*, 38(11):39–41.
- George A. Miller, Claudia Leacock, Randee Tengi, and Ross T. Bunker. 1993. [A semantic concordance](#). In *Proceedings of the workshop on Human Language Technology, HLT '93*, pages 303–308, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Andrea Moro and Roberto Navigli. 2015. [Semeval-2015 task 13: Multilingual all-words sense disambiguation and entity linking](#). In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pages 288–297, Denver, Colorado. Association for Computational Linguistics.
- Andrea Moro, Alessandro Raganato, and Roberto Navigli. 2014. Entity linking meets word sense disambiguation: a unified approach. *TACL*, 2:231–244.
- Roberto Navigli, David Jurgens, and Daniele Vannella. 2013. [SemEval-2013 Task 12: Multilingual Word Sense Disambiguation](#). In *Second Joint Conference on Lexical and Computational Semantics (*SEM), Volume 2: Proceedings of the Seventh International Workshop on Semantic Evaluation (SemEval 2013)*, pages 222–231.
- Roberto Navigli, Kenneth C. Litkowski, and Orin Hargraves. 2007. Semeval-2007 task 07: Coarse-grained english all-words task. In *SemEval-2007*, pages 30–35, Prague, Czech Republic.
- Tommaso Pasini and Roberto Navigli. 2017. [Trainomatic: Large-scale supervised word sense disambiguation in multiple languages without manual training data](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 78–88. Association for Computational Linguistics.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. [Glove: Global vectors for word representation](#). In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543.
- Alain Polguère. 2003. *Lexicologie et sémantique lexicale*. Les Presses de l'Université de Montréal.
- Sameer S. Pradhan, Edward Loper, Dmitriy Dligach, and Martha Palmer. 2007. [Semeval-2007 task 17: English lexical sample, srl and all words](#). In *Proceedings of the 4th International Workshop on Semantic Evaluations*, SemEval '07, pages 87–92, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Alessandro Raganato, Jose Camacho-Collados, and Roberto Navigli. 2017a. [Word sense disambiguation: A unified evaluation framework and empirical comparison](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 99–110, Valencia, Spain. Association for Computational Linguistics.

- Alessandro Raganato, Claudio Delli Bovi, and Roberto Navigli. 2017b. [Neural sequence learning models for word sense disambiguation](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1167–1178. Association for Computational Linguistics.
- Benjamin Snyder and Martha Palmer. 2004. [The english all-words task](#). In *Proceedings of SENSEVAL-3, the Third International Workshop on the Evaluation of Systems for the Semantic Analysis of Text*.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. [Dropout: A simple way to prevent neural networks from overfitting](#). *J. Mach. Learn. Res.*, 15(1):1929–1958.
- Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. [Sequence to sequence learning with neural networks](#). In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2, NIPS’14*, pages 3104–3112, Cambridge, MA, USA. MIT Press.
- Kaveh Taghipour and Hwee Tou Ng. 2015. [One million sense-tagged instances for word sense disambiguation and induction](#). In *Proceedings of the Nineteenth Conference on Computational Natural Language Learning*, pages 338–344, Beijing, China. Association for Computational Linguistics.
- Loïc Vial, Benjamin Lecouteux, and Didier Schwab. 2018. [UFSAC: Unification of Sense Annotated Corpora and Tools](#). In *Language Resources and Evaluation Conference (LREC)*, Miyazaki, Japan.
- Dayu Yuan, Julian Richardson, Ryan Doherty, Colin Evans, and Eric Altendorf. 2016. Semi-supervised word sense disambiguation with neural models. In *COLING 2016*.
- Zhi Zhong and Hwee Tou Ng. 2010. [It makes sense: A wide-coverage word sense disambiguation system for free text](#). In *Proceedings of the ACL 2010 System Demonstrations, ACLDemos ’10*, pages 78–83, Stroudsburg, PA, USA. Association for Computational Linguistics.