

Semi Supervised Preposition-Sense Disambiguation using Multilingual Data

Hila Gonen

Department of Computer Science
Bar-Ilan University
hilagonn@gmail.com

Yoav Goldberg

Department of Computer Science
Bar-Ilan University
yoav.goldberg@gmail.com

Abstract

Prepositions are very common and very ambiguous, and understanding their sense is critical for understanding the meaning of the sentence. Supervised corpora for the preposition-sense disambiguation task are small, suggesting a semi-supervised approach to the task. We show that signals from unannotated multilingual data can be used to improve supervised preposition-sense disambiguation. Our approach pre-trains an LSTM encoder for predicting the translation of a preposition, and then incorporates the pre-trained encoder as a component in a supervised classification system, and fine-tunes it for the task. The multilingual signals consistently improve results on two preposition-sense datasets.

1 Introduction

Preposition-sense disambiguation (Litkowski and Hargraves, 2005; Litkowski and Hargraves, 2007; Schneider et al., 2015; Schneider et al., 2016), is the task of assigning a category to a preposition in context (see Section 2.1). Choosing the correct sense of a preposition is crucial for understanding the meaning of the text. This important semantic task is especially challenging from a learning perspective as only little amounts of annotated training data are available for it. Indeed, previous systems (see Sections 2.1.1 and 5.4) make extensive use of the vast and human-curated WordNet lexicon (Miller, 1995) in order to compensate for the small size of the annotated data and obtain good accuracies.

Instead, we propose to deal with the scarcity of annotated data by taking a semi-supervised approach. We rely on the intuition that word ambiguity tends to differ between languages (Dagan et al., 1991), and show that multilingual corpora can provide a good signal for the preposition sense disambiguation task. Multilingual corpora are vast and relatively easy to obtain (Resnik and Smith, 2003; Koehn, 2005; Steinberger et al., 2006), making them appealing candidates for use in a semi-supervised setting.

Our approach (Section 4) is based on representation learning (Bengio et al., 2013), and can also be seen as an instance of multi-task (Caruana, 1997), or transfer learning (Pan and Yang, 2010). First, we train an LSTM-based neural network (Hochreiter and Schmidhuber, 1997) to predict a foreign (say, French) preposition given the context of an English preposition. This trains the network to map contexts of English prepositions to representations that are predictive of corresponding foreign prepositions, which are in turn correlated with preposition senses. The learned mapper, which takes into account large amounts of parallel text, is then incorporated into a monolingual preposition-sense disambiguation system (Section 3) and is fine-tuned based on the small amounts of available supervised data. We show that the multilingual signal is effective for the preposition-sense disambiguation task on two different datasets (Section 5).

This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License. License details:
<http://creativecommons.org/licenses/by-sa/4.0/>

2 Background

2.1 Proposition Sense Disambiguation

Prepositions are very common, very ambiguous and tend to carry different meanings in different contexts. Consider the following 3 sentences: “You should book a room *for* 2 nights”, “*For* some reason, he is not here yet” and “I went there to get a present *for* my mother”. The preposition “for” has 3 different readings in these sentences: in the first sentence it indicates DURATION, in the second it indicates an EXPLANATION, and in the third a BENEFICIARY. The preposition-sense disambiguation task is defined as follows: given a preposition within a sentential context, decide which category it belongs to, or what its role in the sentence is. Choosing the right sense of a preposition is central to understanding the meaning of an utterance (Baldwin et al., 2009).

2.1.1 Previous Work and Available Corpora

The preposition-sense disambiguation task was the focus of the SemEval 2007 shared task (Litkowski and Hargraves, 2007), based on the set of senses defined in The Preposition Project (TPP) (Litkowski and Hargraves, 2005), with three participating systems (Ye and Baldwin, 2007; Yuret, 2007; Popescu et al., 2007). Since then, it was tackled in several additional works (Dahlmeier et al., 2009; Tratz and Hovy, 2009; Hovy et al., 2010; Tratz, 2011; Srikumar and Roth, 2013b), some of which used different preposition sense inventories and corpora, based on subsets of the TPP dictionary. Srikumar and Roth (2013b) modeled semantic relations expressed by prepositions. For this task, they presented a variation of the TPP inventory, by collapsing related preposition senses, so that all senses are shared between all prepositions (Srikumar and Roth, 2013a). Schneider et al (2015) further improve this inventory and define a new annotation scheme.

There are two main datasets for this task: the corpus of the SemEval 2007 shared task (Litkowski and Hargraves, 2007), and the Web-reviews corpus (Schneider et al., 2016):

SemEval 2007 Corpus This corpus covers 34 prepositions with 16,557 training and 8096 test sentences, each containing a single preposition example. The sentences were extracted from the FrameNet database,¹ based mostly on the British National Corpus (with 75%/25% of informative-writings/literary). Each preposition has a different set of possible senses, with a range of 2 to 25 possible senses for a given preposition. We use the original split to train and test sets.

Web-reviews Corpus Schneider et al (2015) introduce a new, unified and improved sense inventory and corpus (Schneider et al., 2016) in which all prepositions share the same set of senses (senses from a unified inventory are often referred to as supersenses). This corpus contains text in the online reviews genre. It is much smaller than the SemEval corpus, with 4,250 preposition mentions covering 114 different prepositions which are annotated into 63 fine-grained senses. The senses are grouped in a hierarchy, from which we chose a coarse-grained subset of 12 senses for this work: AFFECTOR, ATTRIBUTE, CIRCUMSTANCE, CO-PARTICIPANT, CONFIGURATION, EXPERIENCER, EXPLANATION, MANNER, PLACE, STIMULUS, TEMPORAL, UNDERGOER. We find the Web-reviews corpus more appealing than the SemEval one: the unified sense inventory makes the sense-predictions more suitable for use in downstream applications. While our focus in this work is the Web-reviews corpus, we are the first to report results on this dataset. For the sake of comparison to previous work, we also evaluate our models on the SemEval corpus.

2.2 Neural Networks and Notation

We use $w_{1:n}$ to indicate a list of vectors, and $w_{n:1}$ to indicate the reversed list. We use \circ for vector concatenation, and $x[j]$ for selecting the j^{th} element in a vector x .

A multi-layer perceptron (MLP) is a non linear classifier. In this work, we focus on MLPs with a single hidden layer and a softmax output transformation, and define the function $MLP(x)$ as:

$$MLP(x) = \text{softmax}(U(g(Wx + b1)) + b2)$$

¹<http://framenet.icsi.berkeley.edu/>

where g is a non-linear activation function such as *ReLU* or *tanh*, W and U are input-to-hidden and hidden-to-output transformation matrices, and b_1 and b_2 are optional bias terms. We use subscripts (MLP_{f_1} , MLP_{f_2}) to denote MLPs with different parameters.

Recurrent Neural Networks (RNNs) (Elman, 1990) allow the representation of arbitrary sized sequences, without limiting the length of the history. RNN models have been proven to effectively model sequence-related phenomena such as line lengths, brackets and quotes (Karpathy et al., 2015).

In our implementation we use the long short-term memory network (LSTM), a subtype of the RNN (Hochreiter and Schmidhuber, 1997). $LSTM(w_{1:i})$ is the output vector resulting from inputting the items w_1, \dots, w_i into the LSTM in order.

3 Monolingual Preposition Sense Classification

We start by describing an MLP-based model for classifying prepositions to their senses. For an English sentence $s = w_1, \dots, w_n$ and a preposition position i ,² we classify to the sense y as:

$$y = \underset{j}{\operatorname{argmax}} MLP_{sense}(\phi(s, i))[j]$$

where $\phi(s, i)$ is a feature vector composed of 19 features. The features are based on the features of Tratz and Hovy (2009), and are similar in spirit to those used in previous attempts at preposition sense disambiguation. We deliberately do not include WordNet based features, as we want to focus on features that do not require extensive human-curated resources. This makes our model applicable for use in other languages with minimal change. We use the following features: (1) The embedding of the preposition. (2) The embeddings of the lemmas of the two words before and after the preposition, of the head of the preposition in the dependency tree, and of the first modifier of the preposition. (3) The embeddings of the POS tags of these words, of the preposition, and of the head’s head. (4) The embeddings of the labels of the edges to the head of the preposition, to the head’s head and to the first modifier of the preposition. (5) A boolean that indicates whether one of the two words that follow the preposition is capitalized. The English sentences were parsed using the spaCy parser.³

The network (including the embedding vectors) is trained using cross entropy loss. This model performs relatively well, achieving an accuracy of 73.34 on the Web-reviews corpus, way above the most-frequent-sense baseline of 62.37. On the SemEval corpus, it achieves an accuracy of 74.8, outperforming all participants in the original shared task (Section 5). However, these results are limited by the small size of both training sets. In what follows, we will improve the model using unannotated data.

4 Semi-Supervised Learning Using Multilingual Data

Our goal is to derive a representation from unannotated data that is predictive of preposition-senses. We suggest using multilingual data, following the intuition that preposition ambiguity usually differs between languages (Dagan et al., 1991). For example, consider the following two sentences, taken from the Europarl parallel corpus (Koehn, 2005): “What action will it take to defuse the crisis and tension *in* the region?”, and “These are only available *in* English, which is totally unacceptable”. In the first sentence, the preposition “in” is translated into the French preposition “dans”, whereas in the second one, it is translated into the French preposition “en”. Thus, a representation that is predictive of the preposition’s translation is likely to be predictive also of its sense.

Learning a representation from a multilingual corpus We train a neural network model to encode the context of an English preposition as a vector, and predict the foreign preposition based on the context vector. The resulting context encodings will then be predictive of the foreign prepositions, and hopefully also of the preposition senses.

We derive a training set of roughly 7.4M instances from the Europarl corpus (Koehn, 2005). Europarl contains sentence-aligned data in 21 languages. We started by using several ones, and ended up with a

²We also support multi-word prepositions in this work. The extension is trivial.

³<https://spacy.io/>

subset of 12 languages⁴ that together constitute a good representation of the different language families available in the corpus. Though adding the other languages is possible, we did not experiment with them. To extract the training set, we first word-align⁵ the sentence-aligned data, and then create a dataset of English sentences where each preposition is matched to its translation in a foreign language. Since the alignment of prepositions is noisier than that of content words, we use a heuristic to improve precision: given a candidate foreign-preposition, we verify that the two words surrounding it are aligned to the two words surrounding the English preposition. Additionally, we filter out, for each English preposition, all foreign prepositions that were aligned to it in less than 5% of the cases.

We then train the context representations according to the following model. For an English sentence $s = w_1, \dots, w_n$, a preposition position i and a target preposition p in language L , we encode the context as a concatenation of two LSTMs, one reading the sentence from the beginning up to but not including the preposition, and the other in reverse:

$$ctx(s, i) = LSTM_f(w_{1:i-1}) \circ LSTM_b(w_{n:i+1})$$

This is similar to a BiLSTM encoder, with the difference that the encoding does not include the preposition w_i but only its context. By ignoring the preposition, we force the model to focus on the context, and help it share information between different prepositions. Indeed, including the preposition in the encoder resulted in better performance in foreign preposition classification, but the resulting representation was not as effective when used for the sense disambiguation task.

The context vector is then fed into a language specific MLP for predicting the target preposition:

$$\hat{p} = \operatorname{argmax}_j MLP_L(ctx(s, i))[j]$$

The context-encoder and the word embeddings are shared across languages, but the MLP classifiers that follow are language specific. By using multiple languages, we learn more robust representations.

The English word embeddings can be initialized randomly, or using pre-trained embedding vectors, as we explore in Section 5.1. The network is trained using cross entropy loss, and the error is back-propagated through the context-encoder and the word embeddings.

Using the representation for sense classification Once the encoder is trained over the multilingual data, we incorporate it in the supervised sense-disambiguation model by concatenating the representation obtained from the context encoder to the feature vector. Concretely, the supervised model now becomes:

$$y = \operatorname{argmax}_j MLP_{sense}(ctx(s, i) \circ \phi(s, i))[j]$$

where $ctx(s, i)$ is the output vector of the context-encoder and $\phi(s, i)$ is the feature vector as before.

The network is trained using cross entropy loss, and the error back-propagates also to the context-encoder and to the word embeddings to maximize the model’s ability to adapt to the preposition-sense disambiguation task. The complete model is depicted in Figure 1.

5 Empirical results

Implementation details The models were implemented using PyCNN.⁶ All models were trained using SGD, shuffling the examples before each of the 5 epochs. When training a sense prediction model, we use early stopping and choose the best performing model on the development set. The sense-prediction MLP uses *ReLU* activation, and foreign preposition MLPs use *tanh*, with no bias terms. Unless noted otherwise, we use randomly initialized embedding vectors. For each experiment, we chose the parameters that maximized the accuracy on the dev set.⁷ The accuracies we report are the average accuracies over 5 different seeds.

⁴Bulgarian, Czech, Danish, German, Greek, Spanish, French, Hungarian, Italian, Polish, Romanian and Swedish.

⁵Word-alignment is done using the `cdec` aligner (Dyer et al., 2010).

⁶<https://github.com/clab/cnn>

⁷In most of the experiments, the best results are achieved when the hidden-layer of the sense-prediction MLPs is of the size 500, and the preposition embedding is of size 200. In some cases, the best results are achieved with different dimensions.

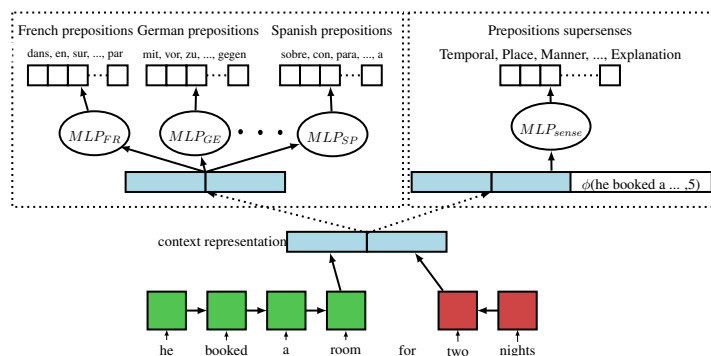


Figure 1: The suggested model for incorporating multilingual data in classifying prepositions to senses. First, a context-encoder (at the bottom, the green and red squares are LSTM cells) is trained on the Europarl corpus, with a different MLP for each language (left dashed frame). Then, the representation obtained from the context-encoder is added to the feature vector when classifying a preposition to senses (right dashed frame).

5.1 Evaluation on the Web-reviews corpus

Using multilingual data Our main motivation in this work was to train a representation which is useful for the preposition-sense disambiguation task. Thus, we compare the performance of our model using the representation obtained from the context-encoder (multilingual model) with the model that does not use this representation (base model). We use the train/test split provided with the corpus. We further split the train set into train and dev sets, by assigning every fourth example of each sense to the dev set, yielding 2552/845/853 instances of train/dev/test.

The results are presented in Table 1. We see an improvement of 2.86 points when using the pre-trained context representations, improving the average result from 73.34 to 76.20.

To verify that the improvement stems from pre-training the context-encoder on multilingual data and not from adding the context-encoder as is, we also evaluated the performance of a model identical to the multilingual model, but with no pre-training on the multilingual data (context model, middle row of Table 1). The context model achieved a very similar result to that of the base model – 73.76, indicating that adding the context-encoder to the base model is not the source of the improvement.

Model	Accuracy
base	73.34 (71.63-73.97)
+context	73.76 (71.86-75.38)
+context(multilingual)	76.20 (74.91-77.26)

Table 1: The average accuracies on the test set of the Web-reviews corpus on 5 different seeds. Numbers in brackets indicate the min and max accuracy across seeds.

Using monolingual or bilingual data only In order to verify the contribution of incorporating information from 12 languages, we also experiment with monolingual and bilingual models. For the monolingual model we train a model similar to our multilingual one, but when trying to predict the English preposition itself, rather than the foreign one, ignoring the multilingual signal altogether. For the bilingual models we train 12 separate models similar to our multilingual model, where each one is trained only on the training examples of a single language.

As shown in Table 2, both the monolingual and the bilingual models improve over the base model (with the exception of Czech), but no improvement is as significant as that of the multilingual model. In addition, we see that the strength of the model does not depend solely on the number of training examples.

Adding external word embeddings Another way of incorporating semi-supervised data into a model is using pre-trained word embeddings. We evaluate our model when using external word embeddings

These two parameters were tuned on the dev set. The embeddings of the features are of dimension 4, with the exception of the lemmas, which are of dimension 50. The dimension of the input to the LSTMs (word embeddings) is 128. Both LSTMs have a single layer with 100 nodes, thus, the representation of the context obtained from the context-encoder is of dimension 200. The hidden-layer of the foreign-preposition MLP is of size 32.

Language	Accuracy	Improvement	Num. of training examples
None (base model)	73.34 (71.63-73.97)	–	–
Czech	73.06 (72.57-73.86)	-0.28	190,850
Polish	73.93 (73.15-74.79)	+0.59	166,101
Italian	73.97 (72.22-75.26)	+0.63	810,589
Romanian	74.09 (73.15-74.56)	+0.75	205,520
Hungarian	74.42 (73.27-75.15)	+1.08	40,302
Bulgarian	74.44 (73.27-74.91)	+1.10	292,908
Spanish	74.65 (73.51-75.73)	+1.31	1,267,400
German	74.73 (73.74-75.62)	+1.39	603,861
Danish	75.08 (74.21-77.49)	+1.74	1,131,915
Greek	75.12 (74.09-76.20)	+1.78	586,494
French	75.43 (74.21-77.02)	+2.09	1,033,267
English (monolingual)	75.68 (74.79-76.55)	+2.34	7,483,206
Swedish	75.87 (74.68-77.49)	+2.53	1,153,999
All 12 languages	76.20 (74.91-77.26)	+2.86	7,483,206

Table 2: The average accuracies on the test set of the Web-reviews corpus on 5 different seeds, using monolingual and bilingual models, along with the improvement over the base model and the number of training examples in each language. Numbers in brackets indicate the min and max accuracy across seeds.

instead of randomly initialized word embeddings. We perform three experiments: 1. using external word embeddings only for the words that are fed into the context-encoder. 2. using external word embeddings only for the lemmas of the features. 3. using external word embeddings for both.

We use two sets of word embeddings: 5-window-bag-of-words-based and dependency-based, both trained by Levy and Goldberg (2014) on English Wikipedia.⁸ As shown in Table 3, both pre-trained embeddings improve the performance of all models in most cases. In all cases, the multilingual model outperforms the base model and the context model, both achieving similar results. Using external word embeddings for both the features and the context-encoder helps the most. The best result of 78.55 is achieved by the multilingual model, improving the result of the base model under the same conditions by 1.71 points.

Model	Context-encoder embeddings only		Feature embeddings only		Embeddings for both	
	Bow	Deps	Bow	Deps	Bow	Deps
base	73.34 (71.63-73.97)	73.34 (71.63-73.97)	76.95 (75.85-77.96)	76.84 (76.32-77.26)	76.95 (75.85-77.96)	76.84 (76.32-77.26)
+context	74.07 (72.10-75.15)	74.42 (73.62-75.03)	76.72 (75.85-77.96)	77.47 (75.85-78.55)	77.14 (76.79-78.08)	77.73 (77.14-78.43)
+context(multilingual)	75.57 (73.51-77.84)	75.90 (75.03-76.55)	77.58 (77.02-78.08)	77.58 (77.14-78.66)	78.45 (77.49-79.48)	78.55 (77.37-79.37)

Table 3: The average accuracies on the test set of the Web-reviews corpus with different pre-trained embeddings on 5 different seeds. Numbers in brackets indicate the min and max accuracy across seeds. **Bow**: 5-words window; **Deps**: dependency-based.

5.2 Evaluation on the SemEval corpus

Adaptations to the SemEval corpus In the SemEval corpus each preposition has a different set of senses, and the natural approach is to learn a different model for each one. We call this the *disjoint* approach. However, we found this approach a bit wasteful in terms of exploiting the annotated data, and we propose a model that uses the information from all prepositions simultaneously (*unified*). In the unified approach, we create an MLP classifier for each preposition, but all of them share a single input-to-hidden transformation matrix and a single bias term. Formally, for a preposition p , we define its MLP as follows:

$$MLP_p(x) = \text{softmax}(U_p(g(Wx + b1)) + b2_p)$$

where W is the shared input-to-hidden transformation matrix, $b1$ is the shared bias term, and U_p and $b2_p$ are preposition-specific hidden-to-output transformation matrix and bias term, respectively. This unified model is trained over the training examples of all prepositions together.

The SemEval corpus sometimes provides multiple senses for a given preposition instance. In both the disjoint and the unified approaches we treat these cases by generalizing the cross entropy loss for multiple correct classes. In the common case, where each training example has a single correct class, the

⁸<https://levyomer.wordpress.com/2014/04/25/dependency-based-word-embeddings/>

cross entropy loss is defined as $-\log p_i$, where p_i is the probability that the model assigns to the correct class. Here, instead of using $-\log p_i$, we use $-\log(\sum_{i \in C} p_i)$, where C is the set of correct classes.

Results The model performs well also on the SemEval corpus, achieving an accuracy of 76.9. Note that we use the exact same parameters that were tuned on the dev set of the Web-reviews corpus, with no additional tuning on this corpus.

As shown in Table 4, the unified model, which trains on all prepositions simultaneously, performs better than a separate model for each preposition (disjoint model), and achieves an improvement of 1.3 points when using the multilingual model. In addition, in both cases we get a significant improvement over the base model when using the pre-trained context-representation. In the unified model, adding the pre-trained context-representation improves the result by 2.1 points. As in the case of the Web-reviews corpus, we can see that this improvement does not stem from adding the context representation as is. Pre-training the representation is essential for achieving these improved results.

Model	Disjoint	Unified
base	73.7 (73.3-74.1)	74.8 (74.4-75.4)
+context	73.8 (73.6-74.0)	75.4 (74.8-75.8)
+context(multilingual)	75.6 (75.4-75.8)	76.9 (76.4-77.7)

Table 4: The average accuracies on the test set of the SemEval corpus on 5 different seeds, with both the disjoint and the unified models. Numbers in brackets indicate the min and max accuracy across seeds.

Similar to the results on the Web-reviews Corpus, when using external word embeddings both for the words that are fed into the context-encoder and for the features, we get an improvement in all models, with an average improvement of 3 points when using the 5-words-window based embeddings. The best result amongst the three models is of 79.6 and is achieved by the multilingual model, improving over the base model by 2.5 points. The results are shown in Table 5.

Note that unlike previous experiments, adding external word embeddings improves the context model over the base model significantly, approaching the results of the multilingual model. For this reason, we also evaluated a model in which we concatenate both contexts: that of the context model (no pre-training), and that of the multilingual model (pre-trained on the multilingual data). In the case where both models achieve similar results, combining both contexts further improves the result, which indicates that they are complementary. The best result of 80.0 is achieved when using both contexts with the 5-window-bag-of-words-based embeddings. We also evaluated this combined model on the Web-reviews corpus, but got no improvement in most cases. This was predictable since in all experiments on that corpus we had a large difference between the results of the context model and of the multilingual model. The only case where we saw an improvement with both contexts was when using dependency-based embeddings for both the features and the context-encoder. The difference between the two datasets can be explained by the much larger size of the SemEval dataset, which allows the context encoder to learn from more data, even without pre-training on multilingual data.

Model	Bow	Deps	None
base	77.1 (76.9-77.2)	76.6 (76.3-76.9)	74.8 (74.4-75.4)
+context	79.5 (78.8-79.9)	78.5 (78.0-78.8)	75.4 (74.8-75.8)
+context(multilingual)	79.6 (79.3-79.9)	79.3 (78.8-79.6)	76.9 (76.4-77.7)
+both contexts	80.0 (79.8-80.2)	79.2 (78.6-79.5)	77.3 (77.2-77.5)

Table 5: The average accuracies on the test set of the SemEval corpus on 5 different seeds, with the unified model, when using external word embeddings for both the context-encoder and the features. Numbers in brackets indicate the min and max accuracy across seeds. **Bow**: 5-words window; **Deps**: dependency-based; **None**: no external word embeddings.

5.3 Using Ensembles

We create an ensemble by training 5 different models (each with a different random seed), and predict test instances using a majority vote over the models. The results are presented in Table 6. As expected, results in all models further improve when using the ensemble. Using the multilingual context helps also when using the ensemble. We see an improvements of 1.99 points on the web-reviews corpus, improving the

result to 80.54. The performance on the SemEval corpus improves by 1.7 points, and reaches an accuracy of 81.7. These results are higher than those of the base model by 2.93 and 2.2 points, respectively.

Model	Web-reviews Corpus		SemEval Corpus	
	Average	Ensemble	Average	Ensemble
base	76.84 (76.32-77.26)	77.61	77.1 (76.9-77.2)	79.5
+context	77.73 (77.14-78.43)	78.90	79.5 (78.8-79.9)	81.1
+context(multilingual)	78.55 (77.37-79.37)	80.54	79.6 (79.3-79.9)	81.2
+both contexts	79.34 (78.43-80.19)	79.84	80.0 (79.8-80.2)	81.7

Table 6: The results on both datasets on 5 different seeds as reported in Tables 3 and 5 in comparison to the results using the ensemble. Numbers in brackets indicate the min and max accuracy across seeds.

5.4 Comparison to previous systems

Table 7 compares our SemEval results with those of previous systems. The system of Ye and Baldwin (2007) got the highest result out of the three participating systems in the SemEval 2007 shared task. They extracted features such as POS tags and WordNet-based features, and also high level features (e.g semantic role tags), using a word window of up to seven words, in a Maximum Entropy classifier. Tratz and Hovy (2009) got a higher result with similar features by using a set of positions that are syntactically related to the preposition instead of a fixed window size. The best performing systems are of Hovy et al (2010) and of Srikumar and Roth (2013b). Both systems rely on vast and thoroughly-engineered feature sets, including many WordNet based features. Hovy et al (2010) explored different word choices (i.e, a fixed window vs. syntactically related words) and different methods of extracting them, while Srikumar and Roth (2013b) improved performance by jointly predicting preposition senses and relations.

In contrast, our models do not include any WordNet based features, making them applicable also for languages lacking such resources. Our models achieve competitive results, outperforming most previous systems, despite using relatively few features and performing hyper-parameter tuning only on the different domain Web-reviews corpus.

Model	Accuracy
base	74.8
+context	75.4
+context(multilingual)	76.9
+context(multilingual) + embeddings	79.6
+both contexts + embeddings	80.0
+both contexts + embeddings + ensemble	81.7
Hovy et al (2010) – using WordNet features	84.8
Srikumar and Roth (2013b) – using WordNet features	84.78
Tratz and Hovy (2009) – using WordNet features	76.4
MELB-YB (Ye and Baldwin, 2007) – using WordNet features	69.3
KU (Yuret, 2007)	54.7
IRST-BP (Popescu et al., 2007)	49.6
Most Frequent Sense	39.6

Table 7: The accuracies on the test set of the SemEval corpus, in comparison to previous systems.

5.5 Error Analysis

Figure 2 depicts the percentage of correct assignments of the base model, in comparison to the multilingual model, per sense and per preposition (only the 10 most common prepositions are shown). Both models use pre-trained word embeddings and ensembles. Clearly, there is a systematic improvement across most prepositions and senses.

6 Related work

Transfer learning and representation learning Transfer learning is a methodology that aims to reduce annotation efforts by first learning a model on a different domain or a closely related task, and then transfer the gained knowledge to the main task (Pan and Yang, 2010). Multi-task learning (MTL) is an approach of transfer learning in which several tasks are trained in parallel while using a shared representation. The different tasks can benefit from each other through this representation (Caruana, 1997). In




Figure 2: Assignments on the dev set of the Web-reviews corpus per preposition (a) and per sense (b). Left bars stand for the base model, right bars stand for the multilingual model. In blue are correct assignments, and in red incorrect ones.

this work we use MTL to improve preposition-sense disambiguation, by using an auxiliary multilingual task – predicting translations of prepositions.

A simple method for sharing information in transfer learning as well as in MTL, is using representations that are shared between related tasks. Representation learning (Bengio et al., 2013) is a closely related field that aims to establish techniques for learning robust and expressive data representations. A well-known effort in this field is that of learning word embeddings for use in a wide range of NLP tasks (Mikolov et al., 2013; Al-Rfou et al., 2013; Levy and Goldberg, 2014; Pennington et al., 2014). While those representations are highly effective in many cases, other scenarios require representations of a full sentence, or of a context around a target word, rather than representations of single words. Contexts are often represented by some manipulation over the embeddings of their words. Such representations have been successfully used for tasks such as context-sensitive similarity (Huang et al., 2012), word sense disambiguation (Chen et al., 2014) and lexical substitution (Melamud et al., 2015). An alternative approach for context representation is encoding a context of arbitrary length into a single vector using LSTMs. This approach has been proven to outperform the previous attempts in a variety of tasks such as Semantic Role Labeling (Zhou and Xu, 2015), Natural Language Inference (Bowman et al., 2015) and Sentence Completion (Melamud et al., 2016). We follow the LSTM-based approach for context representation.

Learning from multilingual data The use of multilingual data for improving monolingual tasks has a long tradition in NLP, and has been used for target word selection (Dagan et al., 1991); word sense disambiguation (Diab and Resnik, 2002); and syntactic parsing and named entity recognition (Burkett et al., 2010), to name a few examples. A dominant approach for exploiting multilingual data is that of cross-lingual projection. This approach assumes a good model exists in one language, and uses annotations in that language in order to constrain possible annotations in another. Projections were successfully used for dependency grammar induction (Ganchev et al., 2009), and for transferring tools such as morphological analyzers and part-of-speech taggers from English to languages with fewer resources (Yarowsky et al., 2001; Yarowsky and Ngai, 2001). A different approach is applying multilingual constraints on existing monolingual models, as done for parsing (Smith and Smith, 2004; Burkett and Klein, 2008) and for morphological segmentation (Snyder and Barzilay, 2008).

Of much relevance to this work are also previous attempts to improve monolingual representations using bilingual data (Faruqui and Dyer, 2014). Previous works focus on creating sense-specific word embeddings instead of the common word-form specific embeddings (Ettinger et al., 2016; Šuster et al., 2016), and also on representing words using their context (Kawakami and Dyer, 2015; Hermann and Blunsom, 2013). While we rely on the assumption most of these works have in common, according to which translations may serve as a strong signal for different senses of words, the novelty of our work is in focusing on prepositions rather than content words, and in jointly representing a context for both a multilingual and a monolingual tasks, which results in an improvement of the monolingual model.

7 Conclusions and Future Work

We show that multilingual data can be used to improve the accuracy of preposition-sense disambiguation. The key idea is to train a context-encoder on vast amounts of parallel data, and by that, to obtain a context representation that is predictive of the sense. We show an improvement of the accuracy in all experiments upon using this representation. Our model achieves an accuracy of 80.54 on the Web-reviews corpus, and an accuracy of 81.7 on the SemEval corpus, with significant improvements over models that do not use the multilingual signals. Our result on the SemEval corpus outperforms most previous works, without using any manually curated lexicons.

Acknowledgements

The work is supported by The Israeli Science Foundation (grant number 1555/15).

References

- Rami Al-Rfou, Bryan Perozzi, and Steven Skiena. 2013. Polyglot: Distributed word representations for multilingual nlp. In *Proceedings of CoNLL 2013*.
- Timothy Baldwin, Valia Kordoni, and Aline Villavicencio. 2009. Prepositions in applications: A survey and introduction to the special issue. *Computational Linguistics*, 35(2):119–149.
- Yoshua Bengio, Aaron Courville, and Pascal Vincent. 2013. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828.
- Samuel R Bowman, Gabor Angeli, Christopher Potts, and Christopher D Manning. 2015. A large annotated corpus for learning natural language inference. In *Proceedings of EMNLP*.
- David Burkett and Dan Klein. 2008. Two languages are better than one (for syntactic parsing). In *Proceedings of EMNLP*, pages 877–886.
- David Burkett, Slav Petrov, John Blitzer, and Dan Klein. 2010. Learning better monolingual models with unannotated bilingual text. In *Proceedings of CoNLL*, pages 46–54.
- Rich Caruana. 1997. Multitask learning. *Machine Learning*, 28:41–75.
- Xinxiong Chen, Zhiyuan Liu, and Maosong Sun. 2014. A unified model for word sense representation and disambiguation. In *Proceedings of EMNLP*, pages 1025–1035.
- Ido Dagan, Alon Itai, and Ulrike Schwall. 1991. Two languages are more informative than one. In *Proceedings of ACL*, pages 130–137.
- Daniel Dahlmeier, Hwee Tou Ng, and Tanja Schultz. 2009. Joint learning of preposition senses and semantic roles of prepositional phrases. In *Proceedings of EMNLP*, pages 450–458.
- Mona Diab and Philip Resnik. 2002. An unsupervised method for word sense tagging using parallel corpora. In *Proceedings of ACL*, pages 255–262.
- Chris Dyer, Adam Lopez, Juri Ganitkevitch, Johnathan Weese, Ferhan Ture, Phil Blunsom, Hendra Setiawan, Vladimir Eidelman, and Philip Resnik. 2010. cdec: A decoder, alignment, and learning framework for finite-state and context-free translation models. In *Proceedings of ACL*.
- Jeffrey L. Elman. 1990. Finding structure in time. *Cognitive science*, 14(2):179–211.
- Allyson Ettinger, Philip Resnik, and Marine Carpuat. 2016. Retrofitting sense-specific word vectors using parallel text. In *Proceedings of NAACL-HLT*, pages 1378–1383.
- Manaal Faruqui and Chris Dyer. 2014. Improving vector space word representations using multilingual correlation. In *Proceedings of EACL*.
- Kuzman Ganchev, Jennifer Gillenwater, and Ben Taskar. 2009. Dependency grammar induction via bitext projection constraints. In *Proceedings of ACL-IJCNLP*, pages 369–377.
- Karl Moritz Hermann and Phil Blunsom. 2013. Multilingual distributed representations without word alignment. In *Proceedings of ICLR*.

- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Dirk Hovy, Stephen Tratz, and Eduard Hovy. 2010. What’s in a preposition? dimensions of sense disambiguation for an interesting word class. In *Proceedings of COLING*.
- 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 ACL*, pages 873–882.
- Andrej Karpathy, Justin Johnson, and Fei-Fei Li. 2015. Visualizing and understanding recurrent networks. *arXiv:1506.02078*.
- Kazuya Kawakami and Chris Dyer. 2015. Learning to represent words in context with multilingual supervision. *arXiv:1511.04623*.
- Philipp Koehn. 2005. Europarl: A parallel corpus for statistical machine translation. In *Proceedings of MT summit*, volume 5, pages 79–86.
- Omer Levy and Yoav Goldberg. 2014. Dependency-based word embeddings. In *Proceedings of ACL*.
- Ken Litkowski and Orin Hargraves. 2005. The preposition project. In *Proceedings of the Second ACL-SIGSEM Workshop on the Linguistic Dimensions of Prepositions and their Use in Computational Linguistics Formalisms and Applications*, pages 171–179.
- Ken Litkowski and Orin Hargraves. 2007. Semeval-2007 task 06: Word-sense disambiguation of prepositions. In *Proceedings of the 4th International Workshop on Semantic Evaluations*, pages 24–29.
- Oren Melamud, Omer Levy, and Ido Dagan. 2015. A simple word embedding model for lexical substitution. In *Proceedings of the 1st Workshop on Vector Space Modeling for Natural Language Processing*, pages 1–7.
- Oren Melamud, Jacob Goldberger, and Ido Dagan. 2016. context2vec: Learning generic context embedding with bidirectional lstm. In *Proceedings of CoNLL*.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*.
- George A Miller. 1995. Wordnet: a lexical database for English. *Communications of the ACM*, 38(11):39–41.
- Sinno Jialin Pan and Qiang Yang. 2010. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10):1345–1359.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of EMNLP 2014*.
- Octavian Popescu, Sara Tonelli, and Emanuele Pianta. 2007. IRST-BP: Preposition disambiguation based on chain clarifying relationships contexts. In *Proceedings of the 4th International Workshop on Semantic Evaluations*.
- Philip Resnik and Noah A Smith. 2003. The web as a parallel corpus. *Computational Linguistics*, 29(3):349–380.
- Nathan Schneider, Vivek Srikumar, Jena D. Hwang, and Martha Palmer. 2015. A hierarchy with, of, and for preposition supersenses. In *Proceedings of the 9th Linguistic Annotation Workshop*, pages 112–123.
- Nathan Schneider, Jena D. Hwang, Vivek Srikumar, Meredith Green, Abhijit Suresh, Kathryn Conger, Tim O’Gorman, and Martha Palmer. 2016. A corpus of preposition supersenses. In *Proceedings of the 10th Linguistic Annotation Workshop*.
- David A Smith and Noah A Smith. 2004. Bilingual parsing with factored estimation: Using english to parse korean. In *Proceedings of EMNLP*.
- Benjamin Snyder and Regina Barzilay. 2008. Cross-lingual propagation for morphological analysis. In *Proceedings of AAAI*.
- Vivek Srikumar and Dan Roth. 2013a. An inventory of preposition relations. *arXiv:1305.5785*.
- Vivek Srikumar and Dan Roth. 2013b. Modeling semantic relations expressed by prepositions. *Transactions of the Association for Computational Linguistics*, 1:231–242.
- Ralf Steinberger, Bruno Pouliquen, Anna Widiger, Camelia Ignat, Tomaz Erjavec, Dan Tufis, and Dániel Varga. 2006. The jrc-acquis: A multilingual aligned parallel corpus with 20+ languages. In *Proceedings of LREC*.

- Simon Šuster, Ivan Titov, and Gertjan van Noord. 2016. Bilingual learning of multi-sense embeddings with discrete autoencoders. *arXiv:1603.09128*.
- Stephen Tratz and Dirk Hovy. 2009. Disambiguation of preposition sense using linguistically motivated features. In *Proceedings of NAACL-HLT Student Research Workshop and Doctoral Consortium*, pages 96–100.
- Stephen Tratz. 2011. *Semantically-enriched parsing for natural language understanding*. Ph.D. thesis, University of Southern California.
- David Yarowsky and Grace Ngai. 2001. Inducing multilingual pos taggers and np bracketers via robust projection across aligned corpora. In *Proceedings of NAACL*.
- David Yarowsky, Grace Ngai, and Richard Wicentowski. 2001. Inducing multilingual text analysis tools via robust projection across aligned corpora. In *Proceedings of HLT*.
- Patrick Ye and Timothy Baldwin. 2007. MELB-YB: Preposition sense disambiguation using rich semantic features. In *Proceedings of the 4th International Workshop on Semantic Evaluations*, pages 241–244.
- Deniz Yuret. 2007. KU: Word sense disambiguation by substitution. In *Proceedings of the 4th International Workshop on Semantic Evaluations*, pages 207–213.
- Jie Zhou and Wei Xu. 2015. End-to-end learning of semantic role labeling using recurrent neural networks. In *Proceedings of ACL*.