

LEARNING TO REPRESENT WORDS IN CONTEXT WITH MULTILINGUAL SUPERVISION

Kazuya Kawakami, Chris Dyer

Department of Computer Science

Carnegie Mellon University

Pittsburgh, PA 15213, USA

{kkawakam, cdyer}@cs.cmu.edu

ABSTRACT

We present a neural network architecture based on bidirectional LSTMs to compute representations of words in the sentential contexts. These context-sensitive word representations are suitable for, e.g., distinguishing different word senses and other context-modulated variations in meaning. To learn the parameters of our model, we use cross-lingual supervision, hypothesizing that a good representation of a word in context will be one that is sufficient for selecting the correct translation into a second language. We evaluate the quality of our representations as features in three downstream tasks: prediction of semantic supersenses (which assign nouns and verbs into a few dozen semantic classes), low resource machine translation, and a lexical substitution task, and obtain state-of-the-art results on all of these.

1 INTRODUCTION

Distributed representations of words, which represent each word as a vector in a low-dimensional space, can be learned from unannotated text corpora using a variety of techniques (Mikolov et al., 2013; Pennington et al., 2014; Landauer & Dumais, 1997). The value of such representations owes to their ability to capture intuitive notions of syntactic and semantic similarity as geometric locality. Despite their empirically proven value as a source of features in many downstream applications (Turian et al., 2010), the “one word type, one vector” assumption made by most word representation models is problematic because words may have multiple meanings.

Two standard solutions to this problem exist. The first to treat each word as a collection of discrete, mutually exclusive senses which are individually represented as vectors (Tian et al., 2014; Neelakantan et al., 2014; Wu & Giles, 2015; Huang et al., 2012; Jauhar et al., 2015). However, identifying the appropriate sense granularity in such models is difficult in practice and in theory (Kilgarriff, 1997; Erk et al., 2013). The second solution, which is the basis of this work, eschews sense inventories (whether latent or explicit) and says that lexical meaning is a function of word and its context (Erk & Padó, 2008; Kintsch, 2001; Mitchell & Lapata, 2008). While previous work has hinted at the promise of this solution, only a small number of hand-crafted word–context composition functions have been considered thus far in the literature on semantic representation learning. This is surprising given the success of learning composition functions for computing phrase and sentence representations (Socher et al., 2011; Kalchbrenner et al., 2014).

There are two central challenges faced by learning to represent words in context. The first is to identifying a suitable function class for the composition function. Such a function must be able to account for the fact that a single word type may have both several completely unrelated meanings as well as a several more or less distinct but still related meanings (Cruse, 2000). For an example of the former, the word *plant* may refer, depending on context, to a factory or to a living organism that photosynthesizes. For an example of the latter, the word *bank* may refer to a financial institution or the building housing a financial institution. Since bidirectional RNN-LSTMs have been shown to be able to learn both compositional (Bahdanau et al., 2014) as well as more arbitrary relationships (Ling et al., 2015), we use these as our composition function class (§2).

The second challenge is to identify an appropriate supervisory signal that will be used to fit the parameters of the function. Our motivating hypothesis—which follows a long line of work in using parallel data as a source of information about semantics (Bannard & Callison-Burch, 2005; Resnik & Yarowsky, 1999; Diab, 2003; Faruqui & Dyer, 2014; Hermann & Blunsom, 2014)—is that a good representation of a word in context will be one that predicts how that word (in its sentential context) translates into a second language (§3). We show that word-in-context representations can be learned efficiently from pairs of words-in-context and single word translations into a second language which are extracted from parallel corpora using a word alignment model.

To evaluate our proposed model and training criterion, we evaluate our learned representations as features in three tasks: supersense tagging, low-resource machine translation (i.e., translation where limited parallel data is available), and a lexical substitution task. Success in each of these requires models that can effectively capturing the meaning of a word in context, and in each, we show our model obtains state-of-the-art performance (§4). Additionally, the feedforward neural net model we use as a baseline for supersense tagging outperforms existing baselines even without our new word-in-context model.

2 MODEL

Our model for contextual words is a bidirectional sequence model based on recurrent neural networks (Chan et al., 2015; Bahdanau et al., 2014; 2015, *inter alia*). Intuitively, this model allows us to condition on arbitrarily long dependencies while having an implicit bias toward more local contexts.

Let $\mathbf{w} = (w_1, w_2, \dots, w_n)$ be the words in a sentence with length n . We also project all words into a fixed d -dimensional vectors $\mathbf{x} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n)$, using a (one-word-per-type) word lookup table.

The model encodes each token of the sentence from left to right according to the standard Long-short term memory recurrences:

$$\begin{aligned} \mathbf{i}_t &= \sigma(\mathbf{W}_{xi}\mathbf{x}_t + \mathbf{W}_{hi}\mathbf{h}_{t-1} + \mathbf{W}_{ci}\mathbf{c}_{t-1} + \mathbf{b}_i) \\ \mathbf{f}_t &= \sigma(\mathbf{W}_{xf}\mathbf{x}_t + \mathbf{W}_{hf}\mathbf{h}_{t-1} + \mathbf{W}_{cf}\mathbf{c}_{t-1} + \mathbf{b}_f) \\ \mathbf{c}_t &= \mathbf{f}_t \odot \mathbf{c}_{t-1} + \mathbf{i}_t \odot \tanh(\mathbf{W}_{xc}\mathbf{x}_t + \mathbf{W}_{hc}\mathbf{h}_{t-1} + \mathbf{b}_c) \\ \mathbf{o}_t &= \sigma(\mathbf{W}_{xo}\mathbf{x}_t + \mathbf{W}_{ho}\mathbf{h}_{t-1} + \mathbf{W}_{co}\mathbf{c}_t + \mathbf{b}_o) \\ \mathbf{h}_t &= \mathbf{o}_t \odot \tanh(\mathbf{c}_t) \end{aligned}$$

This yields a representation $\vec{\mathbf{h}}_t$ for each position in the sentence t which can be interpreted as the representation of word with its left context w_1, w_2, \dots, w_t . The same process is repeated from right to left, yielding a vector $\overleftarrow{\mathbf{h}}_t$. The concatenation of these two vectors

$$\mathbf{h}_t = [\vec{\mathbf{h}}_t; \overleftarrow{\mathbf{h}}_t],$$

is our word-in-context representation.

2.1 MODEL INTUITION

Type-level word embeddings must necessarily represent information about multiple senses in a single vector, and our task is to obtain. To obtain a representation of a word in its context, we want to apply functions which mask or scale some dimensions of the vector according to its context. Thus, functions which apply same scaling function even the word and context are different, average and multi-layer perceptron for example, may not be suitable. The input gate in Long-short term memory is considered to be a suitable scaling function which take target word and its context $(\mathbf{x}_t, \mathbf{h}_t, \mathbf{c}_{t-1})$. Figure 1 show a simplified version of operation to modulate one sense (vegetable plant) from ambiguous type level vector with semantic mask which is conditioned on word and context.

3 MEANING AND TRANSLATION

We now require a training objective that provides supervision for learning the parameters of this model. The question we want to answer is: what is a suitable proxy (or “grounding”) for the mean-




Figure 1: Description of the operation to modulate sense from ambiguous type level vector.

ings of words in context that we can use to construct *token-level* (rather than type-level) word representations?

To illustrate the problem we wish to solve, consider the meaning of the token *bank* in the following sentences:

- I went to the **bank** to deposit my paycheck.
- I went to the river **bank** to eat some lunch.

One very productive strategy for learning semantic word embeddings is to rely on the distributional hypothesis (Harris, 1954), according to which semantically similar items occur in similar contexts. The distributional hypothesis is, furthermore, practically appealing since it enables semantics to be learned from large, unannotated text corpora.

Despite the empirical success of the distributional hypothesis at obtaining representations of word *types*, creating a representation of word *tokens* in terms of context is conceptually unappealing since both the item being embedding and its context potentially share material. One possible solution would be an autoencoding objective, or one might also distinguish between “narrow” and “wide” context (i.e., one that determines the item being embedding and one that provides supervision).¹

However, we instead advocate using an alternative proxy for meaning: how words *translate*. Consider the two examples from above as they might be translated into French.

- Je suis allé à la **banque** pour déposer mon chèque de paie.
- Je suis allé sur la **rive** pour le déjeuner.

The homonymous (i.e., having two completely unrelated senses) word *bank* has been translated into two different words *banque* and *rive* in French.

Finally, while not quite so copious as monolingual corpora, parallel data exist in convenient electronic form in abundance, and this provides a rich resource for learning about the semantics of natural language.

3.1 OBJECTIVE & PARAMETER LEARNING

To operationalize our hypothesis that translation provides a good supervisory signal for learning semantic representations, we learn the parameters of source language word type embeddings and the composition function (i.e., the parameters of the bidirectional LSTMs) by using the computed representation to compute the lexical translation probability of a word in context. That is, we use the computed token embedding to define a probability estimate that a source language word e_t in context $\mathbf{c} = (e_1, \dots, e_{t-1}, e_{t+1}, \dots, e_n)$ translates into a second language as f in vocabulary \mathcal{F} . i.e., $p(f | e_t, \mathbf{c})$.

¹For example, Mikolov et al. (2013) showed that short multiword expressions could be embedding by using the “wider” context that they occur in.

This is done by performing a softmax over the target vocabulary with the representation of the word \mathbf{h}_t , as defined in the previous section. That is, we compute

$$\mathbf{u} = \mathbf{R}\mathbf{h}_t + \mathbf{b}'$$

$$p(f | e_t, \mathbf{c}) = \frac{\exp(u_f)}{\sum_{f' \in \mathcal{F}} \exp(u_{f'})},$$

where parameters \mathbf{R} and \mathbf{b}' define the projection of the source word with context representation \mathbf{h}_t onto the target vocabulary \mathcal{F} .

To obtain pairs of words in context and their lexical translations into a second language, we use unsupervised word alignment techniques (Dyer et al., 2013), to obtain high precision word alignments from a parallel corpus. While modeling alignments as latent variables, or using a soft attention mechanism would be a reasonable alternative, word alignment is fast and the proposed training objective to be easily scaled to large corpora.

Figure 2 illustrates the pre-training architecture.




Figure 2: Description of cross lingual pre-training model.

3.2 PARAMETER LEARNING

The model parameters \mathbf{W} and \mathbf{b} as well as the word projection parameters V_e are first pre-trained with the objective function:

$$\mathcal{L} = - \sum_{(f,e)} \log p(f | e, \mathbf{c})$$

That is, we wish to find the parameters that maximize the lexical translation log probability over the whole parallel corpus of lexical translations (f) of a source word (e) in context (\mathbf{c}).

When we want to transfer the model to another supervised task to predict label $s \in \mathcal{S}$ for a word e in context \mathbf{c} , the final values of the \mathbf{W} and \mathbf{b} parameters are transferred and formulate a similar model to predict label s . Using the transformation matrix $\mathbf{S} \in \mathbb{R}^{|\mathcal{S}| \times d_h}$ and the biases $\mathbf{b}'' \in \mathbb{R}^{|\mathcal{S}|}$, we may define the label probability as

$$\mathbf{u}' = \mathbf{S}\mathbf{h}_t + \mathbf{b}''$$

$$p(s | e_t, \mathbf{c}) = \frac{\exp(u'_s)}{\sum_{s' \in \mathcal{S}} \exp(u'_{s'})}.$$

the model is training by maximizing the log likelihood of the observed label in the task.

$$\mathcal{L}' = - \sum_{(s,e)} \log p(s | e, \mathbf{c}) \tag{1}$$

Table 1: Summary of parallel data.

Dataset	Data Source	Vocabulary		Token		Sentence
		Source	Target	Source	Target	
EN-FR	europarl-v7	93,393	139,934	50,586,497	52,900,470	1,835,733
EN-DE	europarl-v7	93,033	323,367	48,625,466	46,484,368	1,763,744
EN-CS	europarl-v7	51,833	164,770	16,150,983	13,785,699	594,158
EN-FI	europarl-v7	91568	630184	48,584,379	34,819,783	1,779,397
EN-MG	CMU	57,668	76,469	1,592,66	2,023,336	80,306
EN-UR	NIST MT08	43,524	44,566	1,055,030	1,180,031	161,173

4 EXPERIMENTS

We now turn to a series of experiments to show the value of learning representations of words in context according to the objective above. Our paradigm will be to pre-train using the objective above the parameters of a word-in-context model, and then use these (without further fine tuning) in downstream tasks: prediction of semantic supersenses (which assign nouns and verbs into a few dozen semantic classes), low resource machine translation, and a lexical substitution task.

4.1 MODEL CONFIGURATION AND PRE-TRAINING

To pre-train our model, we extracted words (e) in contexts and translations (f) from the Europarl parallel corpus (Koehn, 2005). We conducted experiments with the following four languages: French (FR), German (DE), Czech (CS), Finnish (FI) which are quite typologically diverse. Table 1 shows the numbers of parallel sentences and the numbers of words. For each language pairs, we used 2000 sentences for development and the rest were used for training.

After normal tokenization, we obtained alignments with fast-align tool (Dyer et al., 2013). Since we are modeling single word translations and want high-quality training instances, we run the alignment model in both directions and obtained symmetric alignments by taking intersection between forwards and backward alignments. To control the size of vocabulary, we took 30,000 most common words. For target languages, we removed 10 most common words. The words not in the vocabularies are replaced with $\langle \text{unk} \rangle$ token. We used sentences which have more than 10 words in a sentence.

We used 300 dimension embeddings for source language, and bi-directional LSTMs have 300 hidden units. The trained parameters are source embedding, weights and bias in the model.

We randomly initialized source word embeddings sampled from uniform distribution from -0.08 to 0.08 . All recurrent matrices with orthogonal initialization (Saxe et al., 2013), and non-recurrent weights are initialized from scaled uniform distribution (Glorot & Bengio, 2010). Mini-batches of size 128 are used. We used Adam algorithm for optimization (Kingma & Ba, 2014). We trained models with early-stopping. The perplexities on development data for English to French, German, Czech and Finnish are 3.80, 6.49, 6.30, 19.25 respectively.

4.2 SUPERSENSE TAGGING

Supersenses can be thought of a generalization of words senses into a universal inventory of semantic types. That is, as the number of word senses tend to be too numerous for existing models to generalize properly with the small amounts of data available, supersenses address this problem by clustering all senses into a tractable set of tags. Table 4.2 show examples of supersense tags and its definition. As such, these are generally used in semantically oriented downstream tasks such as co-reference resolution (O’Connor & Heilman, 2013) and question answering (Pasca & Harabagiu, 2001).

Following previous work, we trained our supersense tagger for nouns and verbs on the Sencor dataset. The Sencor datasets consists of three parts, `brown1`, `brown2`, and `brownv`. We mixed these three parts and trained supersense tagger on randomly split 4/5 of data and the rest were used as a development set. We evaluated our model on the held-out SensEval-3 all-words task (Mihalcea

et al., 2004), as done in previous work on supersense tagging (Ciaramita & Altun, 2006; Yuret & Yatbaz, 2010). Since some tokens are annotated with two labels in ambiguous cases, we followed the heuristics of only using the first sense in the data as the correct synset/supersense (Ciaramita & Altun, 2006). To extract supersenses from the Semcor data, we used WordNet version 2.0 synsets.

To avoid the computational overhead of reading extremely wide contexts, we used sliding window to delimit the range of contexts as in (Collobert et al., 2011), that is, each token w_t is embedded using a context window of words $w_{t-n/2}, \dots, w_t, \dots, w_{t+n/2}$. The window size n was fixed to 20.

We use the pre-trained parameters and we put a new task specific softmax layer on top of the hidden units (Fig.2). We updated all parameters including pre-trained parameters. The weights in the softmax layer were initialized from the scaled uniform distribution (Glorot & Bengio, 2010). Mini-batches of size 128 were used with the Adam update rule (Kingma & Ba, 2014).

Since this task has not previously been studied using neural networks, we also report several novel baselines: (1) multi-layer perceptron model which uses a concatenation of a source word type vector and the average of all word type vectors in its context; (2) a forward-only LSTM model; and (3) a bi-directional LSTM with random initialization (rather than cross-lingual pretraining). For fair comparison in terms of the size of word in context representation, we double the hidden unit size of the forward LSTM model.

Supersense	Nouns denoting	Supersense	Verbs denoting
act	acts or actions	change	size, temperature change
artifact	man-made objects	communication	telling, asking, ordering, singing
feeling	feelings and emotions	possession	buying, selling, owning
group	groupings of people or objects	plant	plants
location	spatial position	social	political and social activities

Table 2: Examples of Noun and Verb Supersenses

4.3 LEXICAL TRANSLATION IN LOW RESOURCE LANGUAGE

We investigate the benefit to transfer cross lingually pre-trained word-in-context representation to translation in low-resource language. Since low-resource languages do not have enough data to adequate estimate translation probabilities, we hope that we can learn more effective mappings with pre-trained word-in-context embeddings (Chahuneau et al., 2013).

We trained lexical translation model, which predict translation of aligned English sentence, for low resource languages, Malagasy and Urdu on top of the pre-trained word in context model. Table 1 shows the numbers of parallel sentences and the number of words. We used a dataset used in (Dou et al., 2014) for Malagasy and the Urdu data we used is a part of NIST MT evaluation in 2008-2012². We used 2000 sentences for development and hold-out test set. We filtered out sentences which have less than 3 words for pre-training and words occur less than 1 time are replaced with $\langle \text{unk} \rangle$ token.

We trained our baseline system with cdec (Dyer et al., 2010) and obtained synchronous context-free grammars rules to translate sentences. We added features, translation probability and log translation probability from our translation model and optimized the parameters of a machine translation system with MIRA, Margin-Infused Relaxed Algorithm (Crammer & Singer, 2003).

4.4 LEXICAL SUBSTITUTION

Lexical substitution is the problem of identifying meaning-preserving substitutes for a target word given a sentential context. The task was introduced in SemEval-2007 (McCarthy & Navigli, 2007) involves both finding the synonyms and disambiguating the context. As such, it is an ideal test case for our representations.

Models are evaluated on their ability to predict the substitutes in the gold standard of the LS-SE test-set. We evaluated our model on best and best-mode task which evaluate the quality of the best predictions. The original task allow to make multiple predictions but we only predict only one

²<https://catalog.ldc.upenn.edu/LDC2010T21>

Table 3: Summary of results for supersense tagging.

Method	Semcor			Senseval3		
	Precision	Recall	F1	Precision	Recall	F1
Random	38.2	43.0	40.4	35.8	42.1	38.7
Baseline	63.9	69.3	66.5	60.1	68.7	64.1
HMM	76.7	70.5	77.7	67.6	73.7	70.5
CRF	80.3	80.2	80.2	-	-	-
MLP	82.0	81.9	81.6	83.6	79.7	81.2
LSTM	82.1	82.6	82.1	83.8	81.9	82.5
bi-LSTM	83.5	84.2	83.6	84.6	82.9	83.3
bi-LSTM (FR)	84.8	85.0	84.7	85.8	82.8	84.0
bi-LSTM (DE)	85.2	85.2	85.0	86.2	82.7	84.1
bi-LSTM (CS)	84.9	85.0	84.7	85.9	82.8	84.1
bi-LSTM (FI)	85.0	85.1	84.8	85.9	82.4	83.9
bi-LSTM (average)	85.0	85.1	84.8	86.0	82.7	84.0

Table 4: Summary of results for Translation in low resource Languages.

Method	MG		UR	
	Perplexity ↓	BLEU ↑	Perplexity ↓	BLEU ↑
bi-LSTM (random init)	16.17	21.7	30.87	21.2
bi-LSTM (FR)	12.80	21.9	26.84	21.7
bi-LSTM (DE)	12.97	21.9	26.38	21.4
bi-LSTM (CS)	13.05	22.0	26.21	21.4
bi-LSTM (FI)	13.07	22.0	25.96	21.5
bi-LSTM (average)	12.97	22.0	26.34	21.5

substitution following (Melamud et al., 2015). This task is challenging, since it requires to find the best substitutes from entire word vocabulary.

The way to make prediction is the following. Given a target word and its context, we infer word in context representation of all possible substitutions. Then take one of the most similar words which have highest cosine similarity with target word in context vector as prediction.

For our experiments, we used a simple word alignment base candidate generation to reduce inference time. For a target word in English, we collect all possible French translations from word alignment and took English words 90% most frequently aligned to the French words as candidates. We used same candidates for all our experiments including baseline for fair comparison.

5 RESULT

5.1 SUPERSENSE TAGGING

Table 3 shows frequency weighted Precision, Recall F1 score³ on Semcor test set and Senseval3 all-words task. Our bidirectional LSTM model (bi-LSTM) outperformed the first sense heuristic baseline, the perceptron trained Hidden Markov Model proposed in (Ciaramita & Altun, 2006). And our new word-in-context pre-training model result in further improvements with all language pairs. The averaged score of 4 cross lingually pre-trained models, as in bi-LSTM (average), shows significant improvements over bi-LSTM. The model pre-trained with German achieved best result F1 84.1 on senseval3. Additionally, the baseline neural network models outperforms existing baselines even without cross lingual supervision.

³it can result in an F-score that is not between precision and recall.

Table 5: Summary of results for Lexical Substitution.

Method	best \uparrow	best mode \uparrow
Base	7.81	13.41
Mult	6.64	10.89
BalMult	8.09	13.41
Add	7.37	12.11
BalAdd	8.14	13.41
Skipgram (baseline)	7.77	13.16
bi-LSTM (FR)	9.54	15.79
bi-LSTM (DE)	10.63	18.09
bi-LSTM (CS)	9.74	16.04
bi-LSTM (FI)	8.51	12.99
bi-LSTM (average)	9.60	15.73

Table 6: Disambiguation with multilingual supervision.

Sentence	Translation candidate for <i>plant</i>
They built a large <i>plant</i> to manufacture automobiles.	usine, installation, plante, centrale
Let's <i>plant</i> flowers in the garden.	plantes, planter, végétal, végétale, cultiver

5.2 LEXICAL TRANSLATION IN LOW RESOURCE LANGUAGE

Table 4 shows results on machine translation in low resource language. We report the averaged BLEU score of 5 runs to avoid optimizer randomness Clark et al. (2011). The result show large improvement on perplexity and consistent improvement on BLEU in all language pairs. The average score of 4 cross lingually trained model improved perplexity by around 3 points and BLEU score by 0.3.

5.3 LEXICAL SUBSTITUTION

Table 5 shows results on lexical substitution task. Since our word-in-context representations are build only on Europarl parallel corpora, the baseline system is Skipgram word embedding trained on English side of EN-FR parallel corpora, which is the largest in the corpus. The Skipgram model which take most similar word as prediction is context in-sensitive baseline. Also we compared our results with various context sensitive models, which take arithmetic mean (as in Add and BalAdd) and a geometrical mean (as in Mult and BalMult) of embeddings, proposed by (Melamud et al., 2015). They trained their baseline embeddings (as in Base) on a two billion word web corpus, ukWaC (Ferraresi et al., 2008).

The model achieved best measures⁴ 10.63, best mode measure 18.90 with German supervision. And the second best result was obtained with Czech. As for comparison with Melamud et al. (2015), we cannot compare score directly since we used different corpus and candidate generation. We should compare performance gain by taking into account context. Their best model (BalAdd) achieved 0.33 performance gain with context where our model achieved 2.9 performance gain on best evaluation.

6 DISCUSSION

We proposed the model to predict lexical translation to build word-in-context representation. Table. 6 shows example of disambiguation with translation model in order of translation probability. The model correctly disambiguate industrial plant (usine in French), and vegetable plant (plantes in French). Figure 3 shows the effect of pre-trained word-in-context representation for downstream tasks. Pre-trained model start from low perplexities at the first update and converged earlier, in two epochs, for low resource machine translation.

⁴evaluation was done by a script provided by the task organizer.

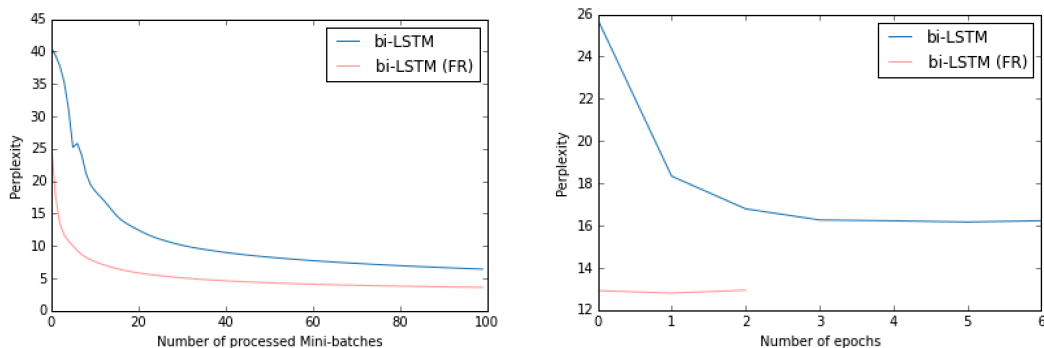


Figure 3: Effect of cross-lingual pre-training for supersense tagging (left) and for low resource machine translation (right). Orange lines are the result from the model pre-trained with French.

We investigated the effect of 4 linguistically diverse language. The results shows the benefit of cross-lingual pre-training in all languages, but overall the model trained with German have stable results and the model trained with Finnish tend to underperform others, especially on lexical substitution task where we do not have supervised fine-tuning process. This is probably because the large vocabulary of Finnish which is two times bigger than German.

7 RELATED WORK

Word representation. Distributed word representations were successfully applied to several downstream tasks such as chunking, parsing, sentiment analysis and paraphrase detection. Most of the tasks requires to use not only word representation but representation of phrases or documents. In the previous works, many architectures were proposed to learn and use word representation. In the sequence modeling problems such as BIO chunking, conditional random fields and recurrent neural networks are applied to represent a sequence of word representations (Turian et al., 2010; Mesnil et al., 2013). For classification tasks such as document classification, sentiment analysis, paraphrase detection, summation of word embeddings (Lauzy et al., 2014), convolutional neural networks (Kalchbrenner et al., 2014) and recursive networks (Socher et al., 2013; Cheng & Kartsaklis, 2015) were proposed to represent compositionality function of words.

Learning semantics from parallel data. Previous works show methods to improve word or document level representation by incorporating multilingual context. Faruqui & Dyer (2014) proposed canonical correlation analysis (CCA) based method to improve the quality of type level representation by projecting word representations of translation pairs (obtained by automatic word alignments) to be maximally correlated in common vector space. Hermann & Blunsom (2013) propose compositional vector space model (CVM) to build sentence representation. They represent a sentence as the sum of its word representations and they train word representation by constraining the representations of parallel sentences to be close. Coulmance et al. (2015) shows that predicting context in target language is an effective way to train word representation shared across languages. Hill et al. (2014) investigated the quality of word embedding learned by neural machine translation model and show its benefit on tasks that require modeling word similarity.

Compositional vector models. Most prior work on compositional vector models has looked primarily at the problem of computing representations of complete phrases rather than specifically words in context. Furthermore, one can learn reasonable generalizations from models that condition on and the generate text using an autoencoding objective (Socher et al., 2011). Dhillon et al. (2012) make the intriguing proposition that left- and right- contexts can be used to supervise each other.

REFERENCES

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.
- Dzmitry Bahdanau, Jan Chorowski, Dmitriy Serdyuk, Philémon Brakel, and Yoshua Bengio. End-to-end attention-based large vocabulary speech recognition. *CoRR*, abs/1508.04395, 2015.
- Colin Bannard and Chris Callison-Burch. Proc. acl. In *Paraphrasing with bilingual parallel corpora*, 2005.
- Victor Chahuneau, Eva Schlinger, Noah A. Smith, and Chris Dyer. Translating into morphologically rich languages with synthetic phrases. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pp. 1677–1687, Seattle, Washington, USA, October 2013. Association for Computational Linguistics. URL <http://www.aclweb.org/anthology/D13-1174>.
- William Chan, Navdeep Jaitly, Quoc V. Le, and Oriol Vinyals. Listen, attend, and spell. *CoRR*, abs/1508.01211, 2015.
- Jianpeng Cheng and Dimitri Kartsaklis. Syntax-aware multi-sense word embeddings for deep compositional models of meaning. *arXiv preprint arXiv:1508.02354*, 2015.
- Massimiliano Ciaramita and Yasemin Altun. Broad-coverage sense disambiguation and information extraction with a supersense sequence tagger. In *Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing*, pp. 594–602. Association for Computational Linguistics, 2006.
- Jonathan H Clark, Chris Dyer, Alon Lavie, and Noah A Smith. Better hypothesis testing for statistical machine translation: Controlling for optimizer instability. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies: short papers-Volume 2*, pp. 176–181. Association for Computational Linguistics, 2011.
- Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. Natural language processing (almost) from scratch. *The Journal of Machine Learning Research*, 12:2493–2537, 2011.
- Jocelyn Coulmance, Jean-Marc Marty, Guillaume Wenzek, and Amine Benhalloum. Trans-gram, fast cross-lingual word-embeddings. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pp. 1109–1113, Lisbon, Portugal, September 2015. Association for Computational Linguistics. URL <http://aclweb.org/anthology/D15-1131>.
- Koby Crammer and Yoram Singer. Ultraconservative online algorithms for multiclass problems. *The Journal of Machine Learning Research*, 3:951–991, 2003.
- Alan Cruse. Meaning in language. *An introduction to Semantics and Pragmatics*, 2000.
- Paramveer S. Dhillon, Jordan Rodu, Dean P. Foster, and Lyle H. Ungar. Two step CCA: A new spectral method for estimating vector models of words. In *Proc. ICML*, 2012.
- Mona Talat Diab. *Word sense disambiguation within a multilingual framework*. PhD thesis, University of Maryland, 2003.
- Qing Dou, Ashish Vaswani, and Kevin Knight. Beyond parallel data: Joint word alignment and decipherment improves machine translation. In *Proceedings of EMNLP*, volume 2014, 2014.
- Chris Dyer, Jonathan Weese, Hendra Setiawan, Adam Lopez, Ferhan Ture, Vladimir Eidelman, Juri Ganitkevitch, Phil Blunsom, and Philip Resnik. cdec: A decoder, alignment, and learning framework for finite-state and context-free translation models. In *Proceedings of the ACL 2010 System Demonstrations*, pp. 7–12. Association for Computational Linguistics, 2010.
- Chris Dyer, Victor Chahuneau, and Noah A. Smith. A simple, fast, and effective reparameterization of IBM model 2. In *Proc. NAACL*, 2013.

- Katrin Erk and Sebastian Padó. A structured vector space model for word meaning in context. In *Proc. EMNLP*, 2008.
- Katrin Erk, Diana McCarthy, and Nicholas Gaylord. Measuring word meaning in context. *Computational Linguistics*, 39(3):511–554, 2013.
- Manaal Faruqui and Chris Dyer. Improving vector space word representations using multilingual correlation. In *Proceedings of EACL*, volume 2014, 2014.
- Adriano Ferraresi, Eros Zanchetta, Marco Baroni, and Silvia Bernardini. Introducing and evaluating ukwac, a very large web-derived corpus of english. In *Proceedings of the 4th Web as Corpus Workshop (WAC-4) Can we beat Google*, pp. 47–54, 2008.
- Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *International conference on artificial intelligence and statistics*, pp. 249–256, 2010.
- Zellig Harris. Distributional structure. *Word*, 10(23):146–162, 1954.
- Karl Moritz Hermann and Phil Blunsom. Multilingual distributed representations without word alignment. *arXiv preprint arXiv:1312.6173*, 2013.
- Karl Moritz Hermann and Phil Blunsom. Multilingual models for compositional distributed semantics. In *Proc. ACL*, 2014.
- Felix Hill, Kyunghyun Cho, Sébastien Jean, Coline Devin, and Yoshua Bengio. Embedding word similarity with neural machine translation. *CoRR*, abs/1412.6448, 2014. URL <http://arxiv.org/abs/1412.6448>.
- Eric H. Huang, Richard Socher, Christopher D. Manning, and Andrew Y. Ng. Improving word representations via global context and multiple word prototypes. In *Proc. ACL*, 2012.
- Sujay K. Jauhar, Chris Dyer, and Eduard Hovy. Ontologically grounded multi-sense representation learning for semantic vector space models. In *Proc. NAACL*, 2015.
- Nal Kalchbrenner, Edward Grefenstette, and Phil Blunsom. A convolutional neural network for modelling sentences. In *Proc. ACL*, 2014.
- Adam Kilgarriff. I don’t believe in word senses. *Computers and the Humanities*, 31(2):91–113, 1997.
- Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Walter Kintsch. Predication. *Cognitive Science*, 25:173–202, 2001.
- Philipp Koehn. Europarl: A parallel corpus for statistical machine translation. In *MT summit*, volume 5, pp. 79–86, 2005.
- Thomas K. Landauer and Susan Dumais. A solution to Plato’s problem: the latent semantic analysis theory of acquisition, induction and representation of knowledge. *Psychological Review*, 104(2): 211–240, 1997.
- Stanislas Lauly, Alex Boulanger, and Hugo Larochelle. Learning multilingual word representations using a bag-of-words autoencoder. *arXiv preprint arXiv:1401.1803*, 2014.
- Wang Ling, Tiago Luís, Luís Marujo, Ramón Fernandez Astudillo, Silvio Amir, Chris Dyer, Alan W Black, and Isabel Trancoso. Finding function in form: Compositional character models for open vocabulary word representation. *arXiv preprint arXiv:1508.02096*, 2015.
- Diana McCarthy and Roberto Navigli. Semeval-2007 task 10: English lexical substitution task. In *Proceedings of the 4th International Workshop on Semantic Evaluations*, pp. 48–53. Association for Computational Linguistics, 2007.

- Oren Melamud, Omer Levy, and Ido Dagan. A simple word embedding model for lexical substitution. In *Proceedings of the 1st Workshop on Vector Space Modeling for Natural Language Processing*, 2015.
- Grégoire Mesnil, Xiaodong He, Li Deng, and Yoshua Bengio. Investigation of recurrent-neural-network architectures and learning methods for spoken language understanding. In *INTER-SPEECH*, pp. 3771–3775, 2013.
- R. Mihalcea, T. Chklovski, and A. Kilgarriff. The senseval-3 english lexical sample task. In *Proceedings of SENSEVAL-3: Third International Workshop on the Evaluation of Systems for the Semantic Analysis of Text*, pp. 25–28, 2004.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. *CoRR*, abs/1301.3781, 2013.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. Distributed representations of words and phrases and their compositionality. In *Proc. NIPS*, 2013.
- Jeff Mitchell and Mirella Lapata. Vector-based models of semantic composition. In *Proc. ACL*, 2008.
- Arvind Neelakantan, Jeevan Shankar, Alexandre Passos, and Andrew McCallum. Efficient non-parametric estimation of multiple embeddings per word in vector space. In *Proc. EMNLP*, 2014.
- Brendan O’Connor and Michael Heilman. Arkref: a rule-based coreference resolution system. *CoRR*, abs/1310.1975, 2013. URL <http://arxiv.org/abs/1310.1975>.
- Marius Pasca and Sanda Harabagiu. The informative role of wordnet in open-domain question answering. In *Proceedings of NAACL-01 Workshop on WordNet and Other Lexical Resources*, pp. 138–143, 2001.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. GloVe: Global vectors for word representation. In *Proc. EMNLP*, 2014.
- Philip Resnik and David Yarowsky. Distinguishing systems and distinguishing senses: new evaluation methods for word sense disambiguation. *Natural Language Engineering*, 5(2):113–133, 1999.
- Andrew M. Saxe, James L. McClelland, and Surya Ganguli. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. *CoRR*, abs/1312.6120, 2013. URL <http://arxiv.org/abs/1312.6120>.
- Richard Socher, Eric H. Huang, Jeffrey Pennington, Andrew Y. Ng, and Christopher D. Manning. Dynamic pooling and unfolding recursive autoencoders for paraphrase detection. In *Proc. NIPS*, 2011.
- Richard Socher, Alex Perelygin, Jean Y Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, and Christopher Potts. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the conference on empirical methods in natural language processing (EMNLP)*, volume 1631, pp. 1642. Citeseer, 2013.
- Fei Tian, Hanjun Dai, Jiang Bian, Bin Gao, Rui Zhang, Enhong Chen, and Tie-Yan Liu. A probabilistic model for learning multi-prototype word embeddings. In *Proc. COLING*, 2014.
- Joseph Turian, Lev Ratinov, and Yoshua Bengio. Word representations: a simple and general method for semi-supervised learning. In *Proc. ACL*, 2010.
- Zhaohui Wu and C. Lee Giles. Sense-aware semantic analysis: A multi-prototype word representation model using Wikipedia. In *Proc. AAAI*, 2015.
- Deniz Yuret and Mehmet Ali Yatzbaz. The noisy channel model for unsupervised word sense disambiguation. *Computational Linguistics*, 36(1):111–127, 2010.