

Learning Bilingual Word Embeddings Using Lexical Definitions

Weijia Shi¹, Muhao Chen¹, Yingtao Tian², Kai-Wei Chang¹

¹Department of Computer Science, University of California, Los Angeles

²Department of Computer Science, Stony Brook University

{swj0419, muhaochen, kw2c}@cs.ucla.edu; {yittian}@cs.stonybrook.edu

Abstract

Bilingual word embeddings, which represent lexicons of different languages in a shared embedding space, are essential for supporting semantic and knowledge transfers in a variety of cross-lingual NLP tasks. Existing approaches to training bilingual word embeddings require often require pre-defined seed lexicons that are expensive to obtain, or parallel sentences that comprise coarse and noisy alignment. In contrast, we propose BilLex that leverages publicly available lexical definitions for bilingual word embedding learning. Without the need of predefined seed lexicons, BilLex comprises a novel word pairing strategy to automatically identify and propagate the precise fine-grained word alignment from lexical definitions. We evaluate BilLex in word-level and sentence-level translation tasks, which seek to find the cross-lingual counterparts of words and sentences respectively. BilLex significantly outperforms previous embedding methods on both tasks.

1 Introduction

Bilingual word embeddings are the essential components of multilingual NLP systems. These embeddings capture cross-lingual semantic transfers of words and phrases from bilingual corpora, and are widely deployed in many NLP tasks, such as machine translation (Conneau et al., 2018), cross-lingual Wikification (Tsai and Roth, 2016), knowledge alignment (Chen et al., 2018) and semantic search (Vulić and Moens, 2015).

A variety of approaches have been proposed to learn bilingual word embeddings (Duong et al., 2017; Luong et al., 2015; Coulmance et al., 2015). Many such approaches rely on the use of aligned corpora. Such corpora could be seed lexicons that provide word-level mappings between two languages (Mikolov et al., 2013a; Xing et al., 2015),

or parallel corpora that align sentences and documents (Klementiev et al., 2012; Gouws et al., 2015). However, these methods critically suffer from several deficiencies. First, seed-lexicon-based approaches are often hindered by the limitedness of seeds (Vulić and Korhonen, 2016), which is an intractable barrier since high-quality seed lexicons require extensive human efforts to obtain (Zhang et al., 2017). Second, parallel corpora provide coarse alignment that does not often accurately infer fine-grained semantic transfers of lexicons (Ruder et al., 2017).

Unlike the existing methods, we propose to use publicly available dictionaries¹ for bilingual word embedding learning. Dictionaries, such as Wiktionary and Merriam-Webster, contain large collections of lexical definitions, which are clean linguistic knowledge that naturally connects word semantics within and across human languages. Hence, dictionaries provide valuable information to bridge the lexicons in different languages. However, cross-lingual learning from lexical definitions is a non-trivial task. A straightforward approach based on aligning the target word embedding to the aggregated embedding of words in the definition might work, but not all words in a definition are semantically related to the defined target word (Fig. 1(a)). Therefore, a successful model has to effectively identify the most related lexicons from the multi-granular and asymmetric alignment of lexical definitions. Besides, how to leverage both bilingual and monolingual dictionaries for cross-lingual learning is another challenge.

In this paper, we propose BilLex (**B**ilingual **W**ord **E**mbeddings **B**ased on **L**exical **D**efinitions) to learn bilingual word embeddings. BilLex constitutes a carefully designed two-stage mechanism

¹We refer to *dictionary* in its regular meaning, i.e. the collections of word definitions. This is different from some papers that refer to dictionaries as seed lexicons.

to automatically cultivate, propagate and leverage lexicon pairs of high semantic similarity from lexical definitions in dictionaries. It first extracts *bilingual strong word pairs* from bilingual lexical definitions of which the words contribute to the cross-lingual definitions of each other. On top of that, our model automatically exploits *induced word pairs*, which utilize monolingual dictionaries and the aforementioned strong pairs to exploit semantically related word pairs. This automated word pair induction process enables BilLex to capture abundant high-quality lexical alignment information, based on which the cross-lingual semantic transfer of words is easily captured in a shared embedding space. Experimental results on word-level and sentence-level translation tasks show that BilLex drastically outperforms various baselines that are trained on parallel or seed-lexicon corpora, as well as state-of-the-art unsupervised methods.

2 Related Work

Prior approaches to learning bilingual word embeddings often rely on word or sentence alignment (Ruder et al., 2017). In particular, seed lexicon methods (Mikolov et al., 2013a; Faruqui and Dyer, 2014; Guo et al., 2015) learn transformations across different language-specific embedding spaces based on predefined word alignment. The performance of these approaches is limited by the sufficiency of seed lexicons. Besides, parallel corpora methods (Gouws et al., 2015; Coulmance et al., 2015) leverage the aligned sentences in different languages and force the representations of corresponding sentence components to be similar. However, aligned sentences merely provide weak alignment of lexicons that do not accurately capture the one-to-one mapping of words, while such a mapping is well-desired by translation tasks (Upadhyay et al., 2016). In addition, a few unsupervised approaches alleviate the use of bilingual resources (Chen and Cardie, 2018; Conneau et al., 2018). These models require considerable effort to train and rely heavily on massive monolingual corpora.

Monolingual lexical definitions have been used for weak supervision of monolingual word similarities (Tissier et al., 2017). Our work demonstrates that dictionary information can be extended to a cross-lingual scenario, for which we develop a simple yet effective induction method to populate fine-grain word alignment.

3 Modeling

We first provide the formal definition of bilingual dictionaries. Let \mathcal{L} be the set of languages and \mathcal{L}^2 be the set of ordered language pairs. For a language $l \in \mathcal{L}$, we use V_l to denote its vocabulary, where for each word $w \in V_l$, $\mathbf{w} \in \mathbb{R}^k$ denotes its embedding vector. A dictionary denoted as D_{l_i, l_j} contains words in language l_i and their definitions in l_j . In particular, D_{l_i, l_j} is a monolingual dictionary if $l_i = l_j$ and is a bilingual dictionary if $l_i \neq l_j$. A dictionary D_{l_i, l_j} contains dictionary entries $(w^i, Q^j(w^i))$, where $w^i \in V_{l_i}$ is the word being defined and $Q^j(w^i)$ is a sequence of words in l_j describing the meaning of the word w^i . Fig. 1(a) shows an entry from an English-French dictionary, and one from a French-English dictionary.

BilLex allows us to exploit semantically related word pairs in two stages. We first use bilingual dictionaries to construct bilingual strong pairs, which are similar to those monolingual word pairs in (Tissier et al., 2017). Then based on the given strong word pairs and monolingual dictionaries, we provide two types of induced word pairs to further enhance the cross-lingual learning.

3.1 Bilingual Strong Pairs

A bilingual strong pair contains two words with high semantic relevance. Such a pair of words that mutually contribute to the cross-lingual definitions of each other is defined as below.

Definition (Bilingual Strong Pairs) P_{l_i, l_j}^S is the set of bilingual strong pairs in $(l_i, l_j) \in \mathcal{L}^2$ ($l_i \neq l_j$), where each word pair is defined as: $(w^i, w^j) \in P_{l_i, l_j}^S \Leftrightarrow w^i \in Q^i(w^j) \wedge w^j \in Q^j(w^i)$

Intuitively, if w^i appears in the cross-lingual definition of w^j and w^j appears in the cross-lingual definition of w^i , then w^i and w^j should be semantically close to each other. Particularly, P_{l_i, l_j}^S denotes monolingual strong pairs if $l_i = l_j$. For instance, $(car, \textit{véhicule})$ depicted in Fig. 1(a) form a bilingual strong pair. Note that Tissier et al. also introduce the monolingual weak pairs by pairing the target word with the other words from its definition, which do not form strong pairs with it. However, we do not extend such weak pairs to the bilingual setting, as we find them to be inaccurate to represent cross-lingual corresponding words.

3.2 Induced Word Pairs

Since bilingual lexical definitions cover only limited numbers of words in two languages, we in-

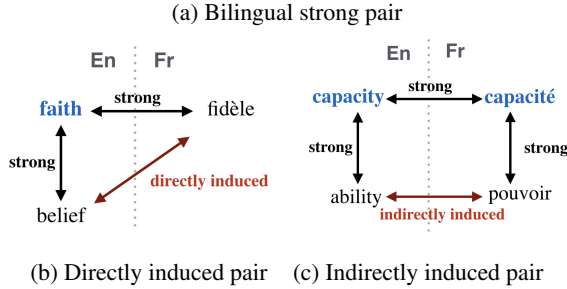
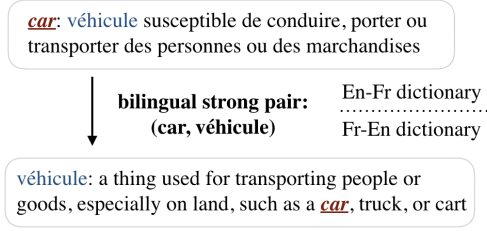


Figure 1: Examples of three types of word pairs. The blue words in (b-c) are pivot words of the induced pairs.

corporate both monolingual and bilingual strong pairs, from which we induce two types of word pairs with different confidence: *directly induced pairs* and *indirectly induced pairs*.

Definition (Bilingual Directly Induced Pairs) P_{l_i, l_j}^D is the set of bilingual directly induced pairs in $(l_i, l_j) \in \mathcal{L}^2$, where each word pair is defined as: $(w^i, w^j) \in P_{l_i, l_j}^D \Leftrightarrow \exists w_p^i, (w^i, w_p^i) \in P_{l_i, l_i}^S \wedge (w_p^i, w^j) \in P_{l_i, l_j}^S$

Intuitively, a bilingual induced pair (w^i, w^j) indicates that we can find a pivot word that forms a monolingual strong pair with one word from (w^i, w^j) and a bilingual strong pair with the other.

Definition (Bilingual Indirectly Induced Pairs) P_{l_i, l_j}^I is the set of bilingual indirectly induced pairs in $(l_i, l_j) \in \mathcal{L}^2$, where each word pair is defined as: $(w^i, w^j) \in P_{l_i, l_j}^I \Leftrightarrow \exists (w_p^i, w_p^j) \in P_{l_i, l_j}^S, (w^i, w_p^i) \in P_{l_i, l_i}^S \wedge (w_p^j, w^j) \in P_{l_j, l_j}^S$

A bilingual indirectly induced pair (w^i, w^j) indicates that there exists a pivot bilingual strong pair (w_p^i, w_p^j) , such that w_p^i forms a monolingual strong pair with w_i and w_p^j forms a monolingual strong pair with w_j . Fig. 1(b-c) shows examples of the two types of induced word pairs.

3.3 Training

Our model jointly learns three word-pair-based cross-lingual objectives Ω_K to align the embedding spaces of two languages, and two monolingual monolingual Skip-Gram losses (Mikolov et al., 2013b) L_{l_i}, L_{l_j} to preserve monolingual word similarities. Given a language pair $(l_i, l_j) \in$

\mathcal{L}^2 , the learning objective of BilLex is to minimize the following joint loss function:

$$J = L_{l_i} + L_{l_j} + \sum_{K \in \{P^S, P^D, P^I\}} \lambda_K \Omega_K$$

Each λ_K ($K \in \{P^S, P^D, P^I\}$) thereof, is the hyperparameter that controls how much the corresponding type of word pairs contributes to cross-lingual learning. For alignment objectives, we use word pairs in both directions of an ordered language pair $(l_i, l_j) \in \mathcal{L}^2$ to capture the cross-lingual semantic similarity of words, such that $P^S = P_{l_i, l_j}^S \cup P_{l_j, l_i}^S$, $P^D = P_{l_i, l_j}^D \cup P_{l_j, l_i}^D$ and $P^I = P_{l_i, l_j}^I \cup P_{l_j, l_i}^I$. Then for each $K \in \{P^S, P^D, P^I\}$, the alignment objective Ω_K is defined as below, where σ is the sigmoid function.

$$\Omega_K = -\frac{1}{|K|} \sum_{(w^i, w^j) \in K} \left(\log \sigma(\mathbf{w}^i \top \mathbf{w}^j) + \sum_{(w_a, w_b) \in N_i(w^j) \cup N_j(w^i)} \log \sigma(-\mathbf{w}_a \top \mathbf{w}_b) \right)$$

For each word pair (w^i, w^j) , we use the unigram distribution raised to the power of 0.75 to select a number of words in l_j (or l_i) for w^i (or w^j) to form a negative sample set $N_i(w^j)$ (or $N_j(w^i)$). Without loss of generality, we define the negative sample set as $N_i(w^j) = \{(w_n^i, w^j) | w_n^i \sim U_i(w) \wedge (w_n^i, w^j) \notin P^S \cup P^D \cup P^I\}$, where $U_i(w)$ is the distribution of words in l_i .

4 Experiment

We evaluate BilLex on two bilingual tasks: *word translation* and *sentence translation retrieval*. Following the convention (Gouws et al., 2015; Mogadala and Rettinger, 2016), we evaluate BilLex between English-French and English-Spanish. Accordingly, we extract word pairs from both directions of bilingual dictionaries in Wiktionary for these language pairs. To support the induced word pairs, we also extract monolingual lexical definitions in the three languages involved, which include 238k entries in English, 107k entries in French and 49k entries in Spanish. The word pair extraction process of BilLex excludes stop words and punctuation in the lexical definitions. The statistics of three types of extracted word pairs are reported in Table 1.

4.1 Word translation

This task aims to retrieve the translation of a source word in the target language. We use the

Lang	#Def	S	I_D	I_I
En&Fr	108,638	52,406	48,524	62,488
En&Es	56,239	32,210	29,857	37,952

Table 1: Statistics of dictionaries and word pair sets.

Language	En-Fr		Fr-En		En-Es		Es-En	
Metric	$P@1$	$P@5$	$P@1$	$P@5$	$P@1$	$P@5$	$P@1$	$P@5$
BiCVM	41.8	56.6	42.6	58.4	37.8	52.8	39.9	54.2
BilBOWA	42.3	59.7	45.2	59.2	37.6	50.3	45.8	53.7
BiSkip	44.0	58.4	45.9	60.2	41.2	58.0	45.4	56.9
Supervised MUSE	74.9	89.8	76.1	90.9	77.4	93.8	77.3	93.6
Unsupervised MUSE	78.1	94.3	78.2	93.0	81.7	94.4	83.3	96.6
BilLex(P^S)	62.4	79.2	61.8	77.4	64.3	78.4	61.9	78.0
BilLex(P^S+P^D)	73.6	87.3	75.3	87.7	73.7	88.7	76.0	90.2
BilLex($P^S+P^D+P^I$)	82.5	96.2	83.8	96.0	82.0	96.5	85.1	96.8

Table 2: Results of the word translation task.

test set provided by [Conneau et al. \(2018\)](#), which selects the most frequent 200k words of each language as candidates for 1.5k query words. We translate a query word by retrieving its k nearest neighbours in the target language, and report $P@k$ ($k = 1, 5$) to represent the fraction of correct translations that are ranked not larger than k .

Evaluation protocol. The hyperparameters of BilLex are tuned based on a small validation set of 1k word pairs provided by [Conneau et al. \(2018\)](#). We allocate 128-dimensional word embeddings with pre-trained BilBOWA ([Gouws et al., 2015](#)), and use the standard configuration to Skip-Gram ([Mikolov et al., 2013b](#)) on monolingual Wikipedia dumps. We set the negative sampling size of bilingual word pairs to 4, which is selected from 0 to 10 with the step of 1. λ_{PS} is set to 0.9, which is tuned from 0 to 1 with the step of 0.1. As we assume that the strong pair relations between words are independent, we empirically set $\lambda_{PD} = (\lambda_{PS})^2 = 0.81$ and $\lambda_{PI} = (\lambda_{PS})^3 = 0.729$. We minimize the loss function using AMSGrad ([Reddi et al., 2018](#)) with a learning rate of 0.001. The training is terminated based on early stopping. We limit the vocabularies as the 200k most frequent words in each language, and exclude the bilingual strong pairs that have appeared in the test set. The baselines we compare against include BiCVM ([Hermann and Blunsom, 2014](#)), BilBOWA ([Gouws et al., 2015](#)), BiSkip ([Luong et al., 2015](#)), supervised and unsupervised MUSE ([Conneau et al., 2018](#)).

Results. Results are summarized in Table 2, where the performance of BilLex is reported for three variants: (i) training with bilingual strong pairs only BilLex(P^S), (ii) with directly induced pair added BilLex(P^S+P^D), and (iii) with all three types of word pairs BilLex($P^S+P^D+P^I$). BilLex($P^S+P^D+P^I$) thereof, offers consistently better performance in all settings, which implies

Language	En-Fr		Fr-En	
Metric	$P@1$	$P@5$	$P@1$	$P@5$
BiCVM	24.4	40.5	32.3	43.8
BilBOWA	27.7	41.4	31.5	47.0
BiSkip	25.3	38.8	26.4	40.4
Supervised MUSE	63.2	76.9	74.9	85.4
Unsupervised MUSE	60.0	76.3	73.7	87.6
BilLex(P^S)	47.4	59.7	57.2	69.6
BilLex(P^S+P^D)	58.7	73.8	67.6	78.9
BilLex($P^S+P^D+P^I$)	64.9	78.2	76.3	89.7

Table 3: Results of sentence translation retrieval.

that the induced word pairs are effective in improving the cross-lingual learning of lexical semantics. Among the baseline models, the unsupervised MUSE outperforms the other four supervised ones. We also discover that for the word translation task, the supervised models with coarse alignment such as BiCVM and BilBOWA do not perform as well as the models with word-level supervision, such as BiSkip and supervised MUSE. Our best BilLex outperforms unsupervised MUSE by 4.4~5.7% of $P@1$ between En and Fr, and by 0.3~1.8% between En and Es. The reason why the settings between En and Fr achieve better performance is that there are much fewer bilingual definitions between En and Es.

4.2 Sentence translation retrieval

This task focuses on retrieving the sentence in the target language space with the tf-idf weighted sentence representation approach. We follow the experiment setup in ([Conneau et al., 2018](#)) with 2k source sentence queries and 200k target sentences from the Europarl corpus for English and French. We carry forward model configurations from the previous experiment, and report $P@k$ ($k = 1, 5$).

Results. The results are reported in Table 3. Overall, our best model variant BilLex($P^S+P^D+P^I$) performs better than the best baseline with a noticeable increment of $P@1$ by 1.4~1.7% and $P@5$ by 1.3~2.1%. This demonstrates that BilLex is suitable for transferring sentential semantics.

5 Conclusion

In this paper, we propose BilLex, a novel bilingual word embedding model based on lexical definitions. BilLex is motivated by the fact that openly available dictionaries offer high-quality linguistic knowledge to connect lexicons across languages. We design the word pair induction method to capture semantically related lexicons in dictionaries, which serve as alignment information in joint training. BilLex outperforms state-of-the-art methods on word and sentence translation tasks.

References

- Muhao Chen, Yingtao Tian, Kai-Wei Chang, Steven Skiena, et al. 2018. Co-training embeddings of knowledge graphs and entity descriptions for cross-lingual entity alignment. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pages 3998–4004.
- Xilun Chen and Claire Cardie. 2018. Unsupervised multilingual word embeddings. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 261–270.
- Alexis Conneau, Guillaume Lample, Marc’Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. 2018. Word translation without parallel data. *ICLR*.
- Jocelyn Coulmance, Jean-Marc Marty, Guillaume Wenzek, and Amine Benhalloum. 2015. Transgram, fast cross-lingual word-embeddings.
- Long Duong, Hiroshi Kanayama, Tengfei Ma, Steven Bird, and Trevor Cohn. 2017. Multilingual training of crosslingual word embeddings. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, volume 1, pages 894–904.
- Manaal Faruqui and Chris Dyer. 2014. Improving vector space word representations using multilingual correlation. In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*, pages 462–471.
- Stephan Gouws, Yoshua Bengio, and Greg Corrado. 2015. Bilbowa: Fast bilingual distributed representations without word alignments. In *International Conference on Machine Learning*.
- Jiang Guo, Wanxiang Che, David Yarowsky, Haifeng Wang, and Ting Liu. 2015. Cross-lingual dependency parsing based on distributed representations. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, volume 1, pages 1234–1244.
- Karl Moritz Hermann and Phil Blunsom. 2014. Multilingual models for compositional distributed semantics. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 58–68.
- Alexandre Klementiev, Ivan Titov, and Binod Bhattarai. 2012. Inducing crosslingual distributed representations of words. In *Proceedings of COLING 2012*.
- Thang Luong, Hieu Pham, and Christopher D Manning. 2015. Bilingual word representations with monolingual quality in mind. In *Proceedings of NAACL-HLT*, pages 151–159.
- Tomas Mikolov, Quoc V Le, and Ilya Sutskever. 2013a. Exploiting similarities among languages for machine translation. *CoRR, abs/1309.4168*.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013b. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119.
- Aditya Mogadala and Achim Rettinger. 2016. Bilingual word embeddings from parallel and non-parallel corpora for cross-language text classification. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 692–702.
- Sashank J Reddi, Satyen Kale, and Sanjiv Kumar. 2018. On the convergence of adam and beyond. In *ICLR*.
- Sebastian Ruder, Ivan Vulić, and Anders Søgaard. 2017. A survey of cross-lingual word embedding models. *Journal of Artificial Intelligence Research*.
- Julien Tissier, Christophe Gravier, and Amaury Habrard. 2017. Dict2vec: Learning word embeddings using lexical dictionaries. In *Conference on Empirical Methods in Natural Language Processing (EMNLP 2017)*, pages 254–263.
- Chen-Tse Tsai and Dan Roth. 2016. Cross-lingual wikification using multilingual embeddings. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 589–598.
- Shyam Upadhyay, Manaal Faruqui, Chris Dyer, and Dan Roth. 2016. Cross-lingual models of word embeddings: An empirical comparison. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 1661–1670.
- Ivan Vulić and Anna Korhonen. 2016. On the role of seed lexicons in learning bilingual word embeddings. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 247–257.
- Ivan Vulić and Marie-Francine Moens. 2015. Monolingual and cross-lingual information retrieval models based on (bilingual) word embeddings. In *Proceedings of the 38th international ACM SIGIR conference on research and development in information retrieval*, pages 363–372. ACM.
- Chao Xing, Dong Wang, Chao Liu, and Yiye Lin. 2015. Normalized word embedding and orthogonal transform for bilingual word translation. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1006–1011.

Meng Zhang, Haoruo Peng, Yang Liu, Huan-Bo Luan, and Maosong Sun. 2017. Bilingual lexicon induction from non-parallel data with minimal supervision. In *AAAI*, pages 3379–3385.