

Weakly-Supervised Concept-based Adversarial Learning for Cross-lingual Word Embeddings

Haozhou Wang¹, James Henderson², Paola Merlo¹

¹ University of Geneva

² Idiap Research Institute

{haozhou.wang, paola.merlo}@unige.ch, james.henderson@idiap.ch

Abstract

Distributed representations of words which map each word to a continuous vector have proven useful in capturing important linguistic information not only in a single language but also across different languages. Current unsupervised adversarial approaches show that it is possible to build a mapping matrix that aligns two sets of monolingual word embeddings without high quality parallel data, such as a dictionary or a sentence-aligned corpus. However, without an additional step of refinement, the preliminary mapping learnt by these methods is unsatisfactory, leading to poor performance for typologically distant languages.

In this paper, we propose a weakly-supervised adversarial training method to overcome this limitation, based on the intuition that mapping across languages is better done at the concept level than at the word level. We propose a concept-based adversarial training method which improves the performance of previous unsupervised adversarial methods for most languages, and especially for typologically distant language pairs.

1 Introduction

Distributed representations of words which map each word to a continuous vector have proven useful in capturing important linguistic information. Vectors of words that are semantically or syntactically similar have been shown to be close to each other in the same space (Mikolov et al., 2013a,c; Pennington et al., 2014), making them widely useful in many natural language processing tasks such as machine translation and parsing (Bansal et al., 2014; Mi et al., 2016), both in a single language and across different languages.

Mikolov et al. (2013b) first observed that the geometric positions of similar words in different languages were related by a linear relation. Zou et al. (2013) showed that a cross-lingually shared word

embedding space is more useful than a monolingual space in an end-to-end machine translation task. However, traditional methods for mapping two monolingual word embeddings require high quality aligned sentences or dictionaries (Faruqui and Dyer, 2014; Ammar et al., 2016b).

Reducing the need for parallel data, then, has become the main issue for cross-lingual word embedding mapping. Some recent work aiming at reducing resources has shown competitive cross-lingual mappings across similar languages, using a pseudo-dictionary, such as identical character strings between two languages (Smith et al., 2017), or a simple list of numerals (Artetxe et al., 2017).

In a more general method, Zhang et al. (2017) have shown that learning mappings across languages via adversarial training (Goodfellow et al., 2014) can avoid using bilingual evidence. This generality comes at the expense of performance. To overcome this limitation, Lample et al. (2018) refine the preliminary mapping matrix trained by generative adversarial networks (GANs) and obtain a model that is again comparable to supervised models for several language pairs. Despite these big improvements, the performance of these refined GAN models depends largely on the quality of the preliminary mappings. It is probably for this reason that these models still do not show satisfactory performance for typologically distant languages.

In this paper, we explore new augmented models of the adversarial framework that are based on the intuition that the right level of lexical alignment between languages, especially distant languages, is not the word but the concept. We also observe that some cross-lingual resources, such as document-aligned data, are not as expensive as high-quality dictionaries. For example, Wikipedia provides thousands of aligned articles in many lan-

guages. Combining these two observations, we develop a concept-based, weakly supervised adversarial method for learning cross-lingual word embedding mappings that rely on these readily available cross-lingual resources.

The main novelty of our method lies in the use of Wikipedia concept-aligned article pairs in the discriminator, which encourages the generator to align words from different languages which are used to describe the same concept. We leverage the power of the adversarial framework and its ability to align distributions by defining concepts as the distribution of words in their respective Wikipedia articles.

The results of our experiments on bilingual lexicon induction show that the preliminary mappings (without post-refinement) trained by our proposed multi-discriminator model are much better than unsupervised adversarial training methods. When we include the post-refinement of Lample et al. (2018), in most cases, our models are comparable or even better than our dictionary-based supervised baselines, and considerably improve the performance of the method of Lample et al. (2018), especially for distant language pairs, such as Chinese-English and Finnish-English.¹ Compared to the previous state-of-the-art, our method still shows competitive performance.

2 Models

The basic idea of mapping two sets of pre-trained monolingual word embeddings together was first proposed by Mikolov et al. (2013b). They use a small dictionary of n pairs of words $\{(w_{s_1}, w_{t_1}), \dots, (w_{s_n}, w_{t_n})\}$ obtained from Google Translate to learn a transformation matrix W that projects the embeddings v_{s_i} of the source language words w_{s_i} onto the embeddings v_{t_i} of their translation words w_{t_i} in the target language, by optimizing the objective:

$$\min_W \sum_{i=1}^n \|Wv_{s_i} - v_{t_i}\|^2 \quad (1)$$

The trained matrix W can then be used for detecting the translation for any source language word w_s by simply finding the word w_t whose embedding v_t is nearest to Wv_s . Recent work by Smith et al. (2017) has shown that forcing the ma-

¹We provide our definition of *similar* and *distant* in section 5.

trix W to be orthogonal can effectively improve the results of mapping.

2.1 GANs for cross-lingual word embeddings

Since the core of this linear mapping is derived from a dictionary, the quality and size of the dictionary can considerably affect the result. Recent attempts by Zhang et al. (2017) and Lample et al. (2018) have shown that, even without a dictionary or any other cross-lingual resource, training the transformation matrix W is still possible using the GAN framework (Goodfellow et al., 2014). The standard GAN framework plays a min-max game between two models: a generative model G and a discriminative model D . The generator G learns from the source data distribution and tries to generate new samples that appear drawn from the distribution of target data. The discriminator D discriminates the generated samples from the target data.

If we adapt the standard GAN framework to the goal of mapping cross-lingual word embeddings, the objective of the generator G is to learn the transformation matrix W that maps the source language embeddings to the target language embedding space, and the discriminator D_l , usually a neural network classifier, detects whether the input is from the target language, giving us the objective:

$$\min_G \max_{D_l} \mathbb{E}_{v_t \sim p_{v_t}} \log D_l(v_t) + \mathbb{E}_{v_s \sim p_{v_s}} \log(1 - D_l(G(v_s))) \quad (2)$$

where $G(v_s) = Wv_s$ and $D_l(v_t)$ denotes the probability that v_t came from the distribution p_{v_t} rather than the distribution p_{v_s} . The inputs of generator G are the embeddings sampled from the distribution of source language word embeddings, p_{v_s} , and the inputs of discriminator D_l are sampled from both the distribution p_{v_t} of real target language word embeddings and the distribution $p_{G(v_s)}$ of generated target language word embeddings. Both G and D_l are trained simultaneously by stochastic gradient descent. A good transformation matrix W can make $p_{G(v_s)}$ similar to p_{v_t} , so that D_l can no longer distinguish between them. However, this kind of similarity is at the distribution level. Good monolingual word embeddings are widely spread in the vector space. Consequently, without any post-processing, the cross-lingual mappings are usually inaccurate when we look at specific words (Lample et al., 2018).

2.2 Concept-based GANs for cross-lingual word embeddings

Under the assumption that words in different languages are similar when they are used to describe the same topic or concept (Søgaard et al., 2015), it seems possible to use concept-aligned documents to improve the mapping performance and make the generated embeddings of the source language words closer to the embeddings of corresponding words in the target language. Unlike dictionaries and sentence-aligned corpora, cross-lingual concept-aligned documents are much more readily available. For example, Wikipedia contains more than 40-million articles in more than 200 languages.² Therefore, many articles of different languages can be linked together because they are about the same concept.

Given two monolingual word embeddings $V_s^j = \{v_{s_1}, \dots, v_{s_j}\}$ and $V_t^k = \{v_{t_1}, \dots, v_{t_k}\}$, our work consists in using a set of Wikipedia concept-aligned article pairs, such that $(C_s^h, C_t^h) = \{(c_{s_1}, c_{t_1}), \dots, (c_{s_h}, c_{t_h})\}$, to reshape the transformation matrix W , where $c_{s_i} \in C_s^h$ represents the article of concept c_i in the source language and $c_{t_i} \in C_t^h$ represents the article of the same concept in the target language.

As shown in Figure 1, the generator uses the transformation matrix W to map the input embeddings v_s from the source to the target language. Differently from the original GANs, the input of the generator and the input of the discriminator on concept D_{cl} are not sampled from the whole distribution of V_s^j or V_t^k , but from their sub-distribution conditioned on the concept c , denoted $p_{v_s|c}$ and $p_{v_t|c}$. v_c represents the embedding of the shared concept c . In this paper, we use the embedding of the title of c in the target language. For titles consisting of multiple words, we average the embeddings of its words.

The input of the discriminator on concept D_{cl} is the concatenation of the embedding of the shared concept v_c and the generated target embeddings $G(v_s)$ or the real target embeddings v_t . Instead of determining whether its input follows the whole target language distribution p_{v_t} , the objective of D_{cl} now becomes to judge whether its input follows the distribution $p_{v_t|c}$, and its loss function is given in equation (3).

$$\log D_{cl}(v_t, v_c) + \log(1 - D_{cl}(G(v_s), v_c)) \quad (3)$$

²https://en.wikipedia.org/wiki/List_of_Wikipedias

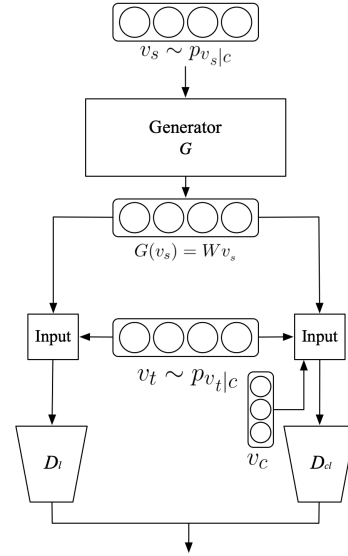


Figure 1: Architecture of our concept-based multi-discriminator model. D_l word-level discriminator, D_{cl} concept-level discriminator.

Because the sub-distribution conditioned on the concept $p_{v_t|c}$ is less spread than p_{v_t} and the proportion of similar words in $p_{v_t|c}$ is higher than in p_{v_t} , the embeddings of the input source words have a greater chance of being trained to align with the embeddings of their similar words in the target language.

Our multi-discriminator concept-based model does not completely remove the usual language discriminator D_l that determines whether its input follows the real target language distribution p_{v_t} or the generated target language distribution, $p_{G(v_s)}$. The simpler discriminator D_l is useful when the concept-based discriminator is not stable. The objective function of the multi-discriminator model, shown in equation (4), combines all these elements.

$$\begin{aligned} \min_G \max_{D_l, D_{cl}} & \mathbb{E}_{v_t \sim p_{v_t}} \log D_l(v_t) \\ & + \mathbb{E}_{v_s \sim p_{v_s}} \log(1 - D_l(G(v_s))) \\ & + \mathbb{E}_{v_t \sim p_{v_t|c}} \log D_{cl}(v_t, v_c) \\ & + \mathbb{E}_{v_s \sim p_{v_s|c}} \log(1 - D_{cl}(G(v_s), v_c)) \end{aligned} \quad (4)$$

3 Shared Components

Both the GAN and the concept-based GAN models use a variety of further methods that have been shown to improve results.

3.1 Orthogonalization

Previous studies (Xing et al., 2015; Smith et al., 2017) show that enforcing the mapping matrix W to be orthogonal can improve the performance and

make the adversarial training more stable. In this work, we perform the same update step proposed by Cisse et al. (2017) to approximate setting W to an orthogonal matrix:

$$W \leftarrow (1 + \beta)W - \beta(WW^T)W \quad (5)$$

According to Lample et al. (2018) and Chen and Cardie (2018), when setting β to less than 0.01, the orthogonalization usually performs well.

3.2 Post-refinement

Previous work has shown that the core cross-lingual mapping can be improved by refining it by bootstrapping from a dictionary extracted from the learned mapping itself (Lample et al., 2018). Since this refinement process is not the focus of this work, we perform the same refinement procedure as Lample et al. (2018). After the core mapping matrix is learned, we use it to translate the ten thousand most frequent source words ($s_1, s_2, \dots, s_{10000}$) into the target language. We then take these ten thousand translations ($t_1, t_2, \dots, t_{10000}$) and translate them back into the source language. We call these translation ($s'_1, s'_2, \dots, s'_{10000}$). We then consider all the word pairs (s_i, t_i) where $s_i = s'_i$ as our seed dictionary and use this dictionary to update our preliminary mapping matrix using the objective in equation 1.

Moreover, following Xing et al. (2015) and Artetxe et al. (2016), we force the refined matrix W^* to be orthogonal by using singular value decomposition (SVD):

$$W^* = ZU^T, U\Sigma Z^T = SVD(V_s^T V_t) \quad (6)$$

3.3 Cross-Domain Similarity Local Scaling

Previous work (Radovanović et al., 2010; Dinu et al., 2015) has shown that standard nearest neighbour techniques are not effective to retrieve target similar words in high-dimensional spaces. The work of Lample et al. (2018) showed that using cross-domain similarity local scaling (CSLS) to retrieve target similar words is more accurate than standard nearest neighbour techniques. Instead of just considering the similarity between the source word and its neighbours in the target language, CSLS also takes into account the similarity between the target word and its neighbours in the source language:

$$CSLS(Wv_s, v_t) = 2 \cos(Wv_s, v_t) - r_t(Wv_s) - r_s(v_t) \quad (7)$$

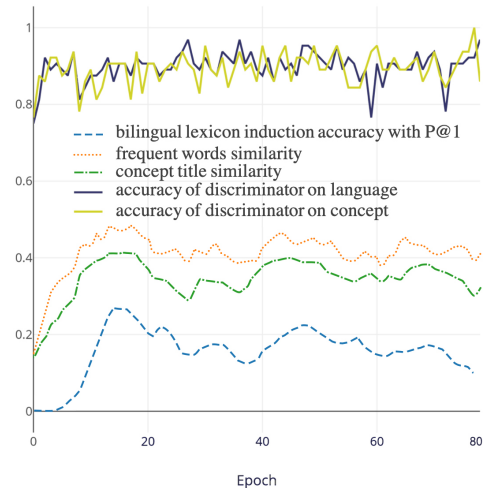


Figure 2: Model selection criteria

where $r_t(Wv_s)$ represents the mean similarity between a source embedding and its neighbours in the target language, and $r_s(v_t)$ represent the mean similarity between a target embedding and its neighbours in the source language. In this work, we use CSLS to build a dictionary for our post-refinement.

4 Training

In our weakly-supervised model, the sampling and model selection procedures are important.

Sampling Procedure As the statistics show in Table 1, even after filtering, the vocabulary of concept-aligned articles is still large. But it has been shown that the embeddings of frequent words are the most informative and useful (Luong et al., 2013; Lample et al., 2018), so we only keep the one-hundred thousand most frequent words for learning W .

For each training step, then, the input word s of our generator is randomly sampled from the vocabulary that is common both to the source monolingual word embedding S and the source Wikipedia concept-aligned articles. After the input source word s is sampled, we sample a concept c according to the frequency of s in each source article of the ensemble of concepts. Then we uniformly sample a target word t from the sub-vocabulary of the target article of concept c .³

Model Selection It is not as difficult to select the best model for weakly-supervised learning as it is

³We use uniform sampling for the target words because the distribution of the sampled sub-vocabulary is very small. Sampling t according to its local frequency in the target article would significantly increase the sampling probability of function words, which have less semantic information.

	Aligned Concepts	Non-English Vocab.	English Vocab.
de \rightleftarrows en	0.16M	1.17M	0.73M
es \rightleftarrows en	0.16M	0.61M	0.69M
fi \rightleftarrows en	0.02M	0.35M	0.17M
fr \rightleftarrows en	0.16M	0.63M	0.71M
it \rightleftarrows en	0.15M	0.56M	0.64M
ru \rightleftarrows en	0.10M	0.89M	0.56M
tr \rightleftarrows en	0.01M	0.24M	0.19M
zh \rightleftarrows en	0.06M	0.39M	0.35M

Table 1: Statistics of Wikipedia concept-aligned articles for our experimentations (M=millions).

for unsupervised learning. In our task, the first selection criterion is the average of the cosine similarity between the embeddings of the concept title in the source and the target languages. Moreover, we find that another good indicator for selecting the best model is the average cosine similarity between the two sets of embeddings consisting of the average embeddings of the ten most frequent words in the source and in the target languages for each aligned article pair. As shown in Figure 2, both of these two criteria correlate well with the word translation evaluation task that we describe in section 5. In this paper, we use the average similarity of the cross-lingual title as the criterion for selecting the best model.

Other Details For our generator, the mapping matrix W is initialized with a random orthogonal matrix. Our two discriminators D_l and D_{cl} are two multi-layer perceptron classifiers with different hidden layer sizes (1024 for D_l and 2048 for D_{cl}), and a ReLU function as the activation function. In this paper, we set the number of hidden layers to 2 for both D_l and D_{cl} . In practice, one hidden layer also performs very well.

5 Experiments

We experimentally evaluate our proposal on the task of Bilingual Lexicon Induction (BLI). This task evaluates directly the bilingual mapping ability of a cross-lingual word embedding model. For each language pair, we retrieve the best translations for source words in the test data, and we report the accuracy.

More precisely, for a given source language word, we map its embedding onto the target language and retrieve the closest word. If this closest word is included in the correct translations in the evaluation dictionary, we consider that it is a correct case. In this paper, we report the results that use CSLS to retrieve translations.

Baselines As the objective of this work is to improve the performance of unsupervised GANs by using our weakly-supervised concept-based model, we choose the model of Lample et al. (2018) (called MUSE below), as our unsupervised baseline, since it is a typical standard GAN model. We evaluate the models both in a setting without refinement and a setting with refinement. The procedure of refinement is described in section 3.2 and 3.3. Moreover, it is important to evaluate whether our model is comparable to previous supervised models. We use two different dictionary-based systems, VecMap proposed by Artetxe et al. (2017) and Procrustes, provided by Lample et al. (2018), as our supervised baselines. Previous experiments have already shown that these two systems are strong baselines.

5.1 Data

Different previous pieces of work on bilingual lexicon induction use different datasets. We choose two from the publicly available ones, selected to have a comprehensive evaluation of our method.

BLI-1 The dataset provided by Lample et al. (2018) contains high quality dictionaries for more than 150 language pairs. For each language pair, it has a training dictionary of 5000 words and an evaluation dictionary of 1500 words. This dataset allows us to have a better understanding of the performance of our method on many different language pairs. We choose nine languages for testing and compare our method to our supervised and unsupervised baselines described in section 5: English (en), German (de), Finnish (fi), French (fr), Spanish (es), Italian (it), Russian (ru), Turkish (tr) and Chinese (zh). We classify *similar* and *distant* languages based on a combination of structural properties (directional dependency distance, as proposed and measured in Chen and Gerdes (2017)) and lexical properties, as measured by the clustering of current large-scale multilingual sentence embeddings.⁴ We consider en \rightleftarrows de, en \rightleftarrows fr, en \rightleftarrows es, en \rightleftarrows it as similar language pairs and en \rightleftarrows fi, en \rightleftarrows ru, en \rightleftarrows tr, en \rightleftarrows zh as distant language pairs.

BLI-2 Unlike BLI-1, the dataset of Dinu et al. (2015) and its extensions provided by Artetxe et al. (2017, 2018b) only consists of dictionaries of 4 language pairs trained on a Europarl parallel cor-

⁴See for example the clustering in <https://code.fb.com/ai-research/laser-multilingual-sentence-embeddings/>

		<i>Supervised</i>		<i>Without refinement</i>		<i>With refinement</i>	
		VecMap	Procrustes	MUSE	Our Method	MUSE	Our Method
Similar language pairs	de-en	72.5	72.0	55.3	<i>66.3</i>	72.2	72.4
	en-de	72.5	72.1	59.2	<i>68.2</i>	72.5	73.4
	es-en	83.4	82.9	<i>78.1</i>	78.0	83.1	83.6
	en-es	81.9	81.5	75.5	74.9	81.1	80.9
	fr-en	81.5	82.2	72.2	<i>75.6</i>	81.2	83.3
	en-fr	81.3	81.3	77.8	<i>78.3</i>	82.2	82.2
	it-en	77.7	79.1	61.0	<i>63.3</i>	76.4	76.2
	en-it	77.2	77.5	63.3	<i>64.3</i>	77.4	78.2
	avg.	78.5	78.6	67.8	<i>71.1</i>	78.3	78.8
Distant language pairs	fi-en	55.9	56.2	0.00	<i>42.8</i>	28.6	53.5
	en-fi	39.5	39.3	0.00	<i>35.2</i>	21.0	43.0
	ru-en	60.5	61.3	43.8	<i>54.3</i>	49.9	58.9
	en-ru	48.0	50.9	28.7	<i>43.2</i>	37.6	51.1
	tr-en	55.7	58.0	12.2	<i>42.1</i>	25.5	54.7
	en-tr	37.3	37.1	26.2	<i>33.1</i>	39.4	41.4
	zh-en	45.0	41.2	26.9	<i>32.1</i>	30.8	42.0
	en-zh	45.4	52.1	29.4	<i>37.7</i>	33.0	48.0
	avg.	48.4	49.5	20.9	<i>40.1</i>	33.2	49.1

Table 2: Results of bilingual lexicon induction (accuracy % P@1) for similar and distant language pairs on the dataset BLI-1. Procrustes and MUSE represent the supervised and unsupervised model of Lample et al. (2018), VecMap is the supervised model of Artetxe et al. (2017). Word translations are retrieved by using CSLs. Bold face indicates the best result overall and italics indicate the best result between the two columns without refinement.

pus. Each dictionary has a training set of 5000 entries and a test set of 1500 entries. Compared to BLI-1, this dataset is much noisier and the entries are selected from different frequency ranges. However, BLI-2 has been widely used for testing by previous methods (Faruqui and Dyer, 2014; Dinu et al., 2015; Xing et al., 2015; Artetxe et al., 2016; Zhang et al., 2016; Artetxe et al., 2017; Smith et al., 2017; Lample et al., 2018; Artetxe et al., 2018a,b).⁵ Using BLI-2 allows us to have a direct comparison with the state-of-the-art.

Monolingual Word Embeddings The quality of monolingual word embeddings has a considerable impact on cross-lingual embeddings (Lample et al., 2018). Compared to CBOW and Skip-gram embeddings, FastText embeddings (Bojanowski et al., 2017) capture syntactic information better. The ideal situation would be to use FastText embeddings for both the BLI-1 and BLI-2 datasets. However, much previous work uses CBOW embeddings, so we use different monolingual word embeddings for BLI-1 and for BLI-2.

For BLI-1, we use FastText⁶ to train our monolingual word embedding models with 300 dimensions for each language and default settings. The

⁵Most of these methods have been tested by Artetxe et al. (2018b) by using their own implementations.

⁶<https://github.com/facebookresearch/fastText>

training corpus comes from a Wikipedia dump.⁷ For European languages, words are lower-cased and tokenized by the scripts of Moses (Koehn et al., 2007).⁸ For Chinese, we first use OpenCC⁹ to convert traditional characters to simplified characters and then use Jieba¹⁰ to perform tokenization. For each language, we only keep the words that appear more than five times.

For BLI-2, following the work of Artetxe et al. (2018a),¹¹ we use their pretrained CBOW embeddings of 300 dimensions. For English, Italian and German, the models are trained on the WacKy corpus. The Finnish model is trained from Common Crawl and the Spanish model is trained from WMT News Crawl.

Concept Aligned Data For concept-aligned articles, we use the Linguatools Wikipedia comparable corpus.¹² The statistics of our concept-aligned data are reported in Table 1.

5.2 Results and Discussion

Table 2 summarizes the results of bilingual lexicon induction on the BLI-1 dataset. We can see

⁷<https://dumps.wikimedia.org/>

⁸<http://www.statmt.org/moses/>

⁹<https://github.com/BYVoid/OpenCC>

¹⁰<https://github.com/fxsjy/jieba>

¹¹In the work of Artetxe et al. (2018a) the authors reimplemented many previous methods described in table 3.

¹²<http://linguatools.org/tools/corpora/wikipedia-comparable-corpora/>

that with post-refinement, our method achieves the best performance on eight language pairs. Table 3 illustrates the results of bilingual lexicon induction on the BLI-2 dataset. Although our method (with refinement) does not achieve the best results, the gap between our method and the best state-of-the-art (Artetxe et al., 2018a) is very small, and in most cases, our method is better than previous supervised methods.

Comparison with unsupervised GANs As we have mentioned before, the preliminary mappings trained by the method of Lample et al. (2018) perform well for some similar language pairs, such as Spanish to English and French to English. After refinement, their unsupervised GANs models can reach the same level as supervised models for these similar language pairs. However, the biggest drawback of standard GANs is that they exhibit poor performance for distant language pairs. The results from Table 2 clearly confirm this. For example, without refinement, the mapping trained by the unsupervised GAN method can only correctly predict 12% of the words from Turkish to English. Given that the quality of preliminary mappings can seriously affect the effect of refinement, the low-quality preliminary mappings for distant language pairs severely limits the improvements brought by post-refinement. Notice that the method of Lample et al. (2018) scores a null result for English to Finnish on both BLI-1 and BLI-2, indicating that totally unsupervised adversarial training can yield rather unpredictable results.

Compared to the method of Lample et al. (2018), the improvement brought by our method is apparent. Our concept-based GAN models perform better than their unsupervised models on almost every language pair for both datasets, BLI-1 and BLI-2. For those languages where no better results are achieved, the gaps from the best are very small. Figure 4 reports the error reduction rate brought by our method for bilingual lexicon induction on the BLI-1 dataset. It can be clearly seen that the improvement brought by our method is more pronounced on distant language pairs than on similar language pairs.

Comparison with supervised baselines Our two dictionary-based supervised baselines are strong. In most cases, the preliminary mappings trained by our concept-based model are not as good, but the gap is small. After post-refinement, our method becomes comparable with these su-

Method	<i>en-de</i>	<i>en-es</i>	<i>en-fi</i>	<i>en-it</i>
Supervised				
Mikolov et al. (2013b)	35.0	27.3	25.9	34.9
Faruqui and Dyer (2014)	37.1	26.8	27.6	38.4
Dinu et al. (2015)	38.9	30.4	29.1	37.7
Xing et al. (2015)	41.3	31.2	28.2	36.9
Artetxe et al. (2016)	41.9	31.4	30.6	39.3
Zhang et al. (2016)	40.8	31.1	28.2	36.7
Artetxe et al. (2017)	40.9	-	28.7	39.7
Smith et al. (2017)	43.3	35.1	29.4	43.1
Artetxe et al. (2018b)	44.1	32.9	44.1	45.3
Unsupervised				
Zhang et al. (2017)	0.00	0.00	0.00	0.00
Lample et al. (2018)	46.8	35.4	0.38	45.2
Artetxe et al. (2018a)	48.2	37.3	32.6	48.1
Weakly-supervised				
Our method	47.7	37.2	30.8	46.4

Table 3: Results of bilingual lexicon induction (accuracy % P@1) on BLI-2 dataset, all the results of previous methods come from (Artetxe et al., 2018a).

performed methods. For some language pairs, such as French-English and for English to Finnish, our method performs better.

Comparison with the state-of-the-art From the results shown in Table 3, we can see that in most cases, our method works better than previous supervised and unsupervised approaches. However, the performance of Artetxe et al. (2018a) is very strong and their method always works better than ours. Two potential reasons may cause this difference: First, their self-learning framework iteratively fine-tunes the transformation until convergence, while our refinement just runs for a certain number of iterations.¹³ Second, their framework consists of many optimization steps, such as symmetric re-weighting of vectors, steps that we do not have.

Mapping examples Table 4 lists some examples of mapping, randomly selected from our experiment results on the dataset of BLI-1. From this table, we can see that for these selected English words, our model performs better than the unsupervised GANs model of Lample et al. (2018), but when we consider the top 3 translation candidates, both models are able to predict the correct translations well. Interestingly, the word "battery" is most commonly translated as Chinese "电池(dianchi)", correctly predicted by our model, and the top 3 translation candidates provided by

¹³As optimisation of the refinement is not the objective of this paper, we follow the work of Lample et al. (2018) and run only five iterations for post-refinement.

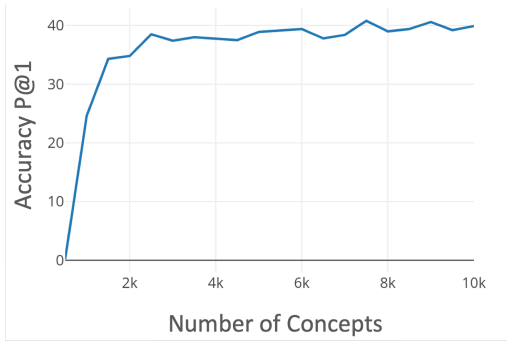


Figure 3: Accuracy of Chinese-English bilingual lexicon induction task for models trained from different concept numbers.




Figure 4: Average error reduction of our method compared to unsupervised adversarial method for bilingual lexicon induction on BLI-1 dataset (Lample et al., 2018). Since the Finnish-English pair is an outlier for the unsupervised method, we report both the average with and without this pair.

our model are all related to "electronic device". However, the translation candidates provided by MUSE are all related to "artillery", a different sense of the the word "battery" both in English and in Chinese, but not as common.

Impact of number of concepts To understand whether the number of aligned concepts affects our method, we trained our concept-based models on a range of Chinese-English concept-aligned article pairs, from 550 to 10'000. We test them on the BLI-1 dataset. Following the trend of performance change shown in Figure 3, we see that, when the number of shared concepts reaches 2500, the improvement in accuracy is already very close to the result of the model trained from the total number of concepts (reported in Table 1), thus indicating a direction for future optimization.

6 Related Work

Sharing a word embedding space across different languages has proven useful for many cross-lingual tasks, such as machine translation (Zou et al., 2013) and cross-lingual dependency pars-

ing (Jiang et al., 2015, 2016; Ammar et al., 2016a). Generally, such spaces can be trained directly from bilingual sentence aligned or document aligned text (Hermann and Blunsom, 2014; Chandar A P et al., 2014; Søgaard et al., 2015; Vulić and Moens, 2013). However the performance of directly trained models is limited by their vocabulary size.

Initial work on the topic has shown that two monolingual spaces can be combined by applying a linear mapping matrix (Mikolov et al., 2013b). This simple approach has been improved upon in several ways: using canonical correlation analysis to map source and target embeddings (Faruqui and Dyer, 2014); or by forcing the mapping matrix to be orthogonal (Xing et al., 2015).

Recently, efforts have concentrated on how to limit or avoid reliance on dictionaries. Good results were achieved with some drastically minimal techniques. Zhang et al. (2016) achieved good results at bilingual POS tagging, but not bilingual lexicon induction, using only ten word pairs to build a coarse orthonormal mapping between source and target monolingual embeddings. The work of Smith et al. (2017) has shown that a singular value decomposition (SVD) method can produce a competitive cross-lingual mapping by using identical character strings across languages. Artetxe et al. (2017, 2018b) proposed a self-learning framework, which iteratively trains its cross-lingual mapping by using dictionaries trained in previous rounds. The initial dictionary of the self-learning can be reduced to 25 word pairs or even only a list of numerals and still have competitive performance. Furthermore, Artetxe et al. (2018a) extend their self-learning framework to unsupervised models, and build the state-of-the-art for bilingual lexicon induction. Instead of using a pre-build dictionary for initialization, they sort the value of the word vectors in both the source and the target distribution, treat two vectors that have similar permutations as possible translations and use them as the initialization dictionary. Additionally, their unsupervised framework also includes many optimization augmentations, such as stochastic dictionary induction and symmetric re-weighting, among others.

Theoretically, employing GANs for training cross-lingual word embedding is also a promising way to avoid the use of bilingual evidence. As far as we know, Miceli Barone (2016) was the first

	MUSE	Our Method
<i>English to Chinese</i>		
film	电影(film), 该片(this film), 本片(this film)	电影(film), 影片(film), 本片(this film)
debate	辩论(debate), 争论(dispute), 讨论(discuss)	辩论(debate), 争论(dispute), 议题(topic)
battery	火炮(artillery), 炮(cannon), 装甲车辆(armored car)	电池(battery), 锂电池(lithium battery), 控制器(controller)
<i>English to French</i>		
january	novembre(november), février(february), janvier(january)	janvier(january) , février(february), septembre(september)
author	auteur(author) , écrivain(writer), romancier(novelist)	auteur(author) , écrivain(writer), romancier(novelist)
polytechnic	université(university), polytechnique(polytechnic) , collège(high school)	université(university), polytechnique(polytechnic) , faculté(faculty)

Table 4: Top 3 English to Chinese and English to French translation candidates selected from our bilingual lexicon induction experiment on the dataset of BLI-1. Word translations are retrieved by using CSLS. Both MUSE and our model are after refinement. Bold face indicates the common translation of selected English words.

attempt at this approach, but the performance of their model is not competitive. Zhang et al. (2017) enforce the mapping matrix to be orthogonal during the adversarial training and achieve good performance on bilingual lexicon induction. The main drawback of their approach is that the vocabularies of their training data are small, and the performance of their models drops significantly when they use large training data. The recent model proposed by Lample et al. (2018) is so far the most successful, becoming competitive with previous supervised approaches through a strong CSLS-based refinement to the core mapping matrix trained by GANs. Even in this case, though, without refinement, the core mappings are not as good as hoped for some distant language pairs. More recently, Chen and Cardie (2018) extends the work of Lample et al. (2018) from the bilingual setting to the multi-lingual setting. Instead of training crosslingual word embeddings for only one language pair, their approach allows them to train crosslingual word embeddings for many language pairs at the same time. Another recent piece of work which is similar to Lample et al. (2018) comes from Xu et al. (2018). Their approach can be divided into 2 steps: first, using Wasserstein GAN (Arjovsky et al., 2017) to train a preliminary mapping between two monolingual distributions and then minimizing the Sinkhorn Distance across distributions. Although their method performs better than Lample et al. (2018) in several

tasks, the improvement mainly comes from the second step, showing that the problem of how to train a better preliminary mapping has not been resolved.

7 Conclusions and Future Work

In this paper, we propose a weakly-supervised adversarial training method for cross-lingual word embedding mapping. Our approach is based on the intuition that mapping across distant languages is better done at the concept level than at the word level. We propose using the distributions of words in aligned Wikipedia articles as data for this concept mapping. The method improves the performance over previous unsupervised adversarial methods in almost all cases, especially for distant languages.

The monolingual word embeddings that we have used in this paper represent a word and not its different senses. These models do not handle polysemous words well. This also can have a negative impact on the performance of cross-lingual mapping. Recent work on contextualized word embeddings (Peters et al., 2018; Devlin et al., 2019) provides dynamic representations for words and has been shown to perform better at handling the problem of polysemy. Extending our concept-based cross-lingual mapping to contextualized word representations will be the focus of our future work.

References

- Waleed Ammar, George Mulcaire, Miguel Ballesteros, Chris Dyer, and Noah A. Smith. 2016a. Many Languages, One Parser. *Transactions of the Association for Computational Linguistics*, 4:431–444.
- Waleed Ammar, George Mulcaire, Yulia Tsvetkov, Guillaume Lample, Chris Dyer, and Noah A. Smith. 2016b. [Massively Multilingual Word Embeddings](#). *arXiv preprint arXiv:1309.4168*.
- Martin Arjovsky, Soumith Chintala, and Léon Bottou. 2017. [Wasserstein GAN](#). *arXiv preprint arXiv:1309.4168*.
- Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2016. Learning principled bilingual mappings of word embeddings while preserving monolingual invariance. In *2016 Conference on Empirical Methods in Natural Language Processing Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2289–2294, Austin, Texas, USA. Association for Computational Linguistics.
- Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2017. Learning bilingual word embeddings with (almost) no bilingual data. In *55th Annual Meeting of the Association for Computational Linguistics*, pages 451–462, Vancouver, Canada.
- Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2018a. A robust self-learning method for fully unsupervised cross-lingual mappings of word embeddings. In *56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 789–798, Melbourne, Australia. Association for Computational Linguistics.
- Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2018b. Generalizing and improving bilingual word embedding mappings with a multi-step framework of linear transformations. In *Thirty-Second Conference on Artificial Intelligence*, pages 5012–5019.
- Mohit Bansal, Kevin Gimpel, and Karen Livescu. 2014. Tailoring Continuous Word Representations for Dependency Parsing. In *52nd Annual Meeting of the Association for Computational Linguistics*, pages 809–815, Baltimore, Maryland, USA.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching Word Vectors with Subword Information. *Transactions of the Association for Computational Linguistics*, 5:135–146.
- Sarath Chandar A P, Stanislas Lauly, Hugo Larochelle, Mitesh Khapra, Balaraman Ravindran, Vikas C Raykar, and Amrita Saha. 2014. An Autoencoder Approach to Learning Bilingual Word Representations. In *2014 Advances in Neural Information Processing Systems*, pages 1853–1861, Montréal Canada.
- Xilun Chen and Claire Cardie. 2018. Unsupervised Multilingual Word Embeddings. In *2018 Conference on Empirical Methods in Natural Language Processing*, pages 261–270. Association for Computational Linguistics.
- Xinying Chen and Kim Gerdes. 2017. [Classifying languages by dependency structure. typologies of delexicalized universal dependency treebanks](#). In *Proceedings of the Fourth International Conference on Dependency Linguistics (Depling 2017)*, pages 54–63, Pisa, Italy. Linköping University Electronic Press.
- Moustapha Cisse, Piotr Bojanowski, Edouard Grave, Yann Dauphin, and Nicolas Usunier. 2017. Parseval Networks: Improving Robustness to Adversarial Examples. In *34th International Conference on Machine Learning*, pages 854–863, Sydney, Australia. PMLR.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Georgiana Dinu, Angeliki Lazaridou, and Marco Baroni. 2015. Improving zero-shot learning by mitigating the hubness problem. In *3rd International Conference on Learning Representations*, pages 1–10, Toulon, France.
- Manaal Faruqui and Chris Dyer. 2014. Improving Vector Space Word Representations Using Multilingual Correlation. In *14th Conference of the European Chapter of the Association for Computational Linguistics*, pages 462–471, Gothenburg, Sweden.
- Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative Adversarial Networks. In *Advances in Neural Information Processing Systems 27*, pages 2672–2680, Montréal, Canada.
- Karl Moritz Hermann and Phil Blunsom. 2014. Multilingual Models for Compositional Distributed Semantics. In *52nd Annual Meeting of the Association for Computational Linguistics*, pages 58–68, Baltimore, Maryland, USA.
- Guo Jiang, Che Wanxiang, Yarowsky David, Wang Haifeng, and Liu Ting. 2015. Cross-lingual dependency parsing based on distributed representations. In *53rd Annual Meeting of the Association for Computational Linguistics and 7th International Joint Conference on Natural Language Processing*, pages 1234–1244, Beijing, China.

- Guo Jiang, Che Wanxiang, Yarowsky David, Wang Haifeng, and Liu Ting. 2016. A distributed representation-based framework for cross-lingual transfer parsing. *Journal of Artificial Intelligence Research*, 55:995–1023.
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondrej Bojar, Alexandra Constantin, and Evan Herbst. 2007. Moses: Open Source Toolkit for Statistical Machine Translation. In *45th Annual Meeting of the Association for Computational Linguistics*, Prague, Czech Republic.
- Guillaume Lample, Alexis Conneau, Marc’Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. 2018. Word translation without parallel data. In *6th International Conference on Learning Representations*, pages 1–14, Vancouver, Canada.
- Minh-Thang Luong, Richard Socher, and Christopher D. Manning. 2013. Better word representations with recursive neural networks for morphology. In *17th Conference on Computational Natural Language Learning*, pages 104–113, Sofia, Bulgaria.
- Haitao Mi, Baskaran Sankaran, Zhiguo Wang, and Abe Ittycheriah. 2016. Coverage Embedding Models for Neural Machine Translation. In *2016 Conference on Empirical Methods in Natural Language Processing*, pages 955–960, Austin, Texas, USA.
- Antonio Valerio Miceli Barone. 2016. Towards cross-lingual distributed representations without parallel text trained with adversarial autoencoders. In *1st Workshop on Representation Learning for NLP*, pages 121–126, Berlin, Germany.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013a. *Efficient Estimation of Word Representations in Vector Space*. *arXiv preprint arXiv:1309.4168*.
- Tomas Mikolov, Quoc V Le, and Ilya Sutskever. 2013b. Exploiting Similarities among Languages for Machine Translation. *arXiv preprint arXiv:1309.4168*, (Computation and Language):1–10.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013c. Distributed Representations of Words and Phrases and their Compositionality. In *26th International Conference on Neural Information Processing Systems*, pages 3111–3119, Lake Tahoe, Nevada, USA.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove - Global Vectors for Word Representation. In *2014 Conference on Empirical Methods in Natural Language Processing*, pages 1532–1543, Doha, Qatar.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep Contextualized Word Representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Miloš Radovanović, Alexandros Nanopoulos, and Mirjana Ivanović. 2010. Hubs in Space: Popular Nearest Neighbors in High-Dimensional Data. *Journal of machine learning research*, 11(Sep):2487–2531.
- Samuel L. Smith, David H. P. Turban, Steven Hamblin, and Nils Y. Hammerla. 2017. Offline bilingual word vectors, orthogonal transformations and the inverted softmax. In *15th International Conference on Learning Representations*, pages 1–10, Toulon, France.
- Anders Søgaard, Zeljko Agić, Héctor Martínez Alonso, Barbara Plank, Bernd Bohnet, and Anders Johannsen. 2015. Inverted indexing for cross-lingual NLP. In *53rd Annual Meeting of the Association for Computational Linguistics and 7th International Joint Conference on Natural Language Processing*, pages 1713–1722, Beijing, China.
- Ivan Vulić and Marie-Francine Moens. 2013. A Study on Bootstrapping Bilingual Vector Spaces from Non-Parallel Data (and Nothing Else). In *2013 Conference on Empirical Methods in Natural Language Processing*, pages 1613–1624, Seattle, Washington, USA.
- Chao Xing, Dong Wang, Chao Liu, and Yiye Lin. 2015. Normalized Word Embedding and Orthogonal Transform for Bilingual Word Translation. In *2015 Annual Conference of the North American Chapter of the ACL*, pages 1006–1011, Denver, Colorado, USA. Association for Computational Linguistics.
- Ruochen Xu, Yiming Yang, Naoki Otani, and Yuexin Wu. 2018. Unsupervised Cross-lingual Transfer of Word Embedding Spaces. In *2018 Conference on Empirical Methods in Natural Language Processing*, pages 2465–2474. Association for Computational Linguistics.
- Meng Zhang, Yang Liu, Huanbo Luan, and Maosong Sun. 2017. Adversarial Training for Unsupervised Bilingual Lexicon Induction. In *55th Annual Meeting of the Association for Computational Linguistics*, pages 1959–1970, Vancouver, Canada. Association for Computational Linguistics.
- Yuan Zhang, David Gaddy, Regina Barzilay, and Tommi Jaakkola. 2016. Ten Pairs to Tag - Multilingual POS Tagging via Coarse Mapping between Embeddings. In *15th Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 12–17, San Diego, California, USA. Association for Computational Linguistics.

Will Y Zou, Richard Socher, Daniel Cer, and Christopher D. Manning. 2013. Bilingual Word Embeddings for Phrase-Based Machine Translation. In *Conference on Empirical Methods in Natural Language Processing*, pages 1393–1398, Seattle, Washington, USA.