

ABSent: Cross-Lingual Sentence Representation Mapping with Bidirectional GANs

Zuohui Fu, Yikun Xian, Shijie Geng, Yingqiang Ge, Yuting Wang, Xin Dong, Guang Wang, Gerard de Melo

Department of Computer Science

Rutgers University, New Brunswick, NJ, USA

{zuohui.fu, sg1309, yingqiang.ge, xd48}@rutgers.edu, siriusxyk@gmail.com, {yw632, gw255}@cs.rutgers.edu, gdm@demelo.org

Abstract

A number of cross-lingual transfer learning approaches based on neural networks have been proposed for the case when large amounts of parallel text are at our disposal. However, in many real-world settings, the size of parallel annotated training data is restricted. Additionally, prior cross-lingual mapping research has mainly focused on the word level. This raises the question of whether such techniques can also be applied to effortlessly obtain cross-lingually aligned sentence representations. To this end, we propose an Adversarial Bidirectional Sentence Embedding Mapping (ABSent) framework, which learns mappings of cross-lingual sentence representations from limited quantities of parallel data. The experiments show that our method outperforms several technically more powerful approaches, especially under challenging low-resource circumstances. The source code is available from <https://github.com/zuohuif/ABSent> along with relevant datasets.

1 Introduction

Not only have regular vector representations of words become ubiquitous (Mikolov et al. 2013), but bilingual word embeddings (Ruder 2017) have as well enjoyed remarkable success, enabling many novel forms of cross-lingual NLP. Owing to the availability of adequate cross-lingual word datasets and supervised learning methods, cross-lingual transfer learning at the word level is well-studied. However, only few results exist on sentence-level cross-lingual mapping, let alone studying low-resource settings.

In order to effortlessly obtain sentence-level representations, one of the most popular methods is to compute a (possibly weighted) average of the word vectors of all words encountered in a sentence. The method is favored for its straightforward simplicity, particularly in light of the widespread availability of pre-trained word vectors. While there are numerous more powerful methods (cf. Section 2), they require substantially longer training times and pre-trained models are typically less convenient to load. Somewhat surprisingly, weighted averages of word vectors have been shown capable of outperforming several more advanced techniques, including certain Long short-term memory (LSTM)-based setups (Wieting et al. 2015;

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: Sentence alignment between the respective embeddings of English and German sentences. Solid circles represent training samples, while hollow circles represent test samples. The challenge is to align test samples (connected by dotted segments) with limited numbers of training pairs (connected by solid segments) and a large number of unpaired training samples.

Arora, Liang, and Ma 2017). In this paper, we study to what extent sentence embeddings based on simple word vector averages can be aligned cross-lingually. While word vector averages have been studied for embeddings of entire text documents, such document embeddings mainly need to capture topic information. Sentence embeddings, in contrast, are typically expected to retain more detailed semantic information. If simple word vector averages can achieve this cross-lingually despite being entirely oblivious of the order of words in the input sentences, this would provide a simple means of connecting semantically related sentences across language boundaries, in support of a diverse range of possible tasks such as question answering, recommendation, plagiarism detection (Liu et al. 2019b; Xian et al. 2019; Ferrero et al. 2017). While there has been research on joint multilingual training of NMT to obtain richer cross-lingual sentence embeddings (Wang et al. 2019), such methods tend to require substantial training data.

At the same time, aligning sentences cross-lingually with limited parallel data is challenging, as it is not obvious how to exploit non-parallel data (see Figure 1). While linear transformations have proven fruitful for cross-lingual word vector mapping (Mikolov, Le, and Sutskever 2013), recent




Figure 2: The framework of our proposed ABSent method. It learns two generators G_X and G_Y to approximate the joint distribution of vectors from both languages. G_X projects sentence embeddings x from language X to Y , while, conversely, G_Y projects sentence embeddings y from language Y to X .

work shows that non-linear transformations may be necessary even just at the level of individual words (Nakashole and Flauger 2018). At the sentence level, this necessity may be much more pronounced, due to the divergent syntactic and morphological properties of different languages and the linear superposition of different concepts. We not only study this empirically but also show how non-linear transformations can be learned with limited parallel data.

Specifically, we propose Adversarial Bi-directional Sentence Embedding Mapping (ABSent), based on the generative adversarial network (GAN) framework to bridge the gap between languages while avoiding overfitting even with limited parallel data. We consider simple (weighted) averaged word embeddings for a source language sentence as input but generates a sentence embedding in a target language space resulting from (weighted) averaged word vectors in a target language. The bi-directional structure additionally enables joint transformations between two languages.

The major contributions of the paper can be outlined as follows: 1) We highlight the simplicity and effectiveness of inducing high-quality sentence representations from pre-trained word embeddings by means of Term Frequency-Inverse Document Frequency (TF-IDF) weighted averaging word vectors. 2) We propose an adversarially bidirectional model for cross-lingual sentence embedding mapping based on a custom form of GAN framework that is capable of utilizing non-parallel sentence pairs. Moreover, we show that the same model architecture can easily be extended to more than one source language. 3) We extensively evaluate the performance of our method on the Tatoeba and Europarl corpora, obtaining exceptional accuracy as well as high quality mapping results, even in low-resource settings.

2 Related Work

Cross-Lingual Projection Approaches. A number of papers consider linear projections to align two word vector spaces with a regression objective (Mikolov, Le, and Sutskever 2013; Zou et al. 2013). Faruqui and Dyer (2014) proposed using Canonical Correlation Analysis (CCA). Xing et al. (2015) showed that adding an orthogonality constraint to the mapping can significantly enhance the result

quality, and has a closed-form solution. There have been approaches that assume that languages share some common vocabulary items as a heuristic for supervision (Smith et al. 2017; Dong and De Melo 2018; Artetxe, Labaka, and Agirre 2017).

A few works also attempt to align monolingual word vector spaces with no supervision at all. Zhang et al. (2017) employed a form of adversarial training, but their approach differs from ours in multiple respects. First, they rely on sharp drops of discriminator accuracy for model selection. Second, their performance is highly sensitive to the selected parallel corpus. Lample et al. (2018) presented a related unsupervised technique that learns a rotation matrix that outperforms several state-of-the-art supervised techniques. In contrast to our approach, none of the above methods consider non-linear transformations (Nakashole and Flauger 2018).

Sentence Embeddings. Well-known approaches to create sentence embeddings include the Paragraph Vector approach (Le and Mikolov 2014), which straightforwardly extends word2vec to generate vectors for paragraphs, and the Skip-Thought Vector approach (Kiros et al. 2015), which relies on recurrent units to encode and decode sentence representations such that these are predictive of neighbouring sentences. There are more sophisticated methods that rely on supervision from a range of different NLP tasks (Subramanian et al. 2018; Yang et al. 2019).

However, inspired by the results from Wieting et al. (2015), Arora, Liang, and Ma (2017) presented a weighting technique that enables simple weighted sums of word vectors to outperform several state-of-the-art models. In our experiments, we build on these insights and as well consider weighted sums of word vectors as sentence embeddings, as these are readily available, even for many low-resource languages. Our weighting scheme is described in Section 4.

Adversarial Training. Recently, GANs (Goodfellow et al. 2014) have shown remarkable success across a diverse range of multimodal tasks. Their adversarial training process resembles a min-max game. Some GAN approaches require a supervised learning setting like image-to-image

transfer (Isola et al. 2017). The CycleGAN approach (Zhu et al. 2017) shows promise in its exploitation of unpaired data to achieve a domain transfer. While GANs have mostly been considered for multimodal data (Liu et al. 2019a), we show how they can be used for linguistic representations in an NLP task.

3 ABSent Approach

In this paper, we seek to learn a transformation between two languages such that the mapping model can be invoked to project an embedding of a source language sentence to a target language space and be able to find the nearest neighbor targets in the target space at the sentence level. At the same time, our approach is shown to be robust under low-resource conditions in terms of the amount of parallel sentences available for training. We start with a formal definition of our sentence representation problem under limited parallel data, followed by a detailed demonstration of our proposed deep neural model.

3.1 Problem Definition

Formally, we assume a source language (domain) X and a target one Y , such that each element is a d -dimensional vector, denoted by $\mathbf{x} \in X$ or $\mathbf{y} \in Y$. We assume that $\forall \mathbf{x} \in X$, \mathbf{x} is aligned with one $\mathbf{y} \in Y$, denoted by $(\mathbf{x}, \mathbf{y}) \in \mathcal{D}$. Given a bilingual corpus \mathcal{D} and a labeled (parallel) subset $\mathcal{D}_l \subseteq \mathcal{D}$, an unlabeled subset \mathcal{D}_{ul} and some distance measure f_d , our goal is to learn two non-linear transformation functions G_X and G_Y that minimize

$$\mathcal{L}_d = \mathbb{E}_{\mathbf{x}, \mathbf{y}} [f_d(\mathbf{x}, G_Y(\mathbf{y})) + f_d(\mathbf{y}, G_X(\mathbf{x}))] \quad (1)$$

Note that in the labeled set \mathcal{D}_l , the alignment between \mathbf{x} is known, while in the unlabeled set \mathcal{D}_{ul} , the relationship between \mathbf{x} and \mathbf{y} are unknown.

In this paper, we consider the case when the labeled set $|\mathcal{D}_l|$ is very limited and the size of the unlabeled set $|\mathcal{D}_{ul}|$ is large. In other words, the challenge is how to utilize unpaired vectors from two domains to learn good mappings from one domain to the other. Specifically, the model is expected to be able to jointly learn from bidirectional transformations between the two languages at the same time to improve the mappings for each direction by better modeling the joint distribution, which is important in settings with limited parallel data as considered in this paper.

3.2 Our Method

To solve the transformation problem with limited parallel data, we introduce the novel *Adversarial Bi-directional Sentence Embedding Mapping* (ABSent) method.

The fundamental core of our model is inspired by the Triangle Generative Adversarial Network (Gan et al. 2017), which addresses the task of image-to-image translation. The key idea is that the generator component learns to non-linearly project embeddings across the two representation spaces, while a discriminator component attempts to distinguish automatically projected embeddings from genuine target language embeddings, thus constraining us to more closely match the target distribution. Unlike regular GANs,

our model incorporates additional information from “adversarial” pairs of sentence embeddings that come from both parallel and non-parallel data. We define the corresponding objective function as follows.

$$\begin{aligned} \mathcal{L}_{\text{real}} = & \mathbb{E}_{\mathbf{x}, \mathbf{y}} [\log(D_{\text{real}}(\mathbf{x}, \mathbf{y}))] \\ & + \mathbb{E}_{\mathbf{x}} [\log(1 - D_{\text{real}}(\mathbf{x}, G_X(\mathbf{x})))] \\ & + \mathbb{E}_{\mathbf{y}} [\log(1 - D_{\text{real}}(\mathbf{y}, G_Y(\mathbf{y})))] \end{aligned} \quad (2)$$

Here, D_{real} is a discriminator that aims to distinguish *real pairs* from *fake pairs*. A *real pair* $(\mathbf{x}, \mathbf{y}) \in \mathcal{D}_l$ is a known mapping in the parallel dataset. A *fake pair* $(\mathbf{x}, G_X(\mathbf{x}))$ or $(\mathbf{y}, G_Y(\mathbf{y}))$ is an artificial pair based on a projection emitted by the generator G_X or G_Y .

Equation 2 reflects an adversarial min-max game, in which the generators G_X , G_Y and the discriminator D_{real} are trained adversarially and concurrently to improve their respective abilities. This is a bidirectional process due to its reliance on both generator functions G_X and G_Y to map from one domain to the other.

In addition, in order to utilize non-parallel information, we further take into consideration *mismatch pairs* induced from non-parallel data. Given the set of source language sentence embeddings X and the set of target language sentence embeddings Y , the set of *mismatch pairs* consists of all training pairs $(\mathbf{x}', \mathbf{y}') \in X \times Y$ such that \mathbf{x}' is an embedding for a sentence that is *not* translationally equivalent to the sentence represented by \mathbf{y}' . The loss function is augmented with the following mismatch term:

$$\mathcal{L}_{\text{mis}} = \mathbb{E}_{\mathbf{x}', \mathbf{y}'} [\log(1 - D_{\text{real}}(\mathbf{x}', \mathbf{y}'))] \quad (3)$$

By combining Equations 2 and 3, we force the discriminator D_{real} to distinguish *real pairs* (\mathbf{x}, \mathbf{y}) from unpaired data, which includes both *mismatch pairs* $(\mathbf{x}', \mathbf{y}')$ as well as generated *fake pairs* $(\mathbf{x}, G_X(\mathbf{x}))$ or $(\mathbf{y}, G_Y(\mathbf{y}))$.

However, the discriminator D_{real} alone cannot determine the directionality between *fake pairs*. Therefore, we introduce another discriminator D_{dom} to distinguish whether fake pairs come from the X domain or from the Y domain. The loss function is defined as:

$$\begin{aligned} \mathcal{L}_{\text{dom}} = & \mathbb{E}_{\mathbf{x}} [\log(D_{\text{dom}}(\mathbf{x}, G_X(\mathbf{x})))] \\ & + \mathbb{E}_{\mathbf{y}} [\log(1 - D_{\text{dom}}(\mathbf{y}, G_Y(\mathbf{y})))] \end{aligned} \quad (4)$$

The overall framework of our ABSent method is illustrated in Figure 2, where we seek to solve the following joint optimization problem:

$$\begin{aligned} \mathcal{L} = & \mathcal{L}_{\text{real}} + \mathcal{L}_{\text{mis}} + \mathcal{L}_{\text{dom}} \\ G_X^*, G_Y^* = & \arg \min_{G_X, G_Y} \max_{D_{\text{real}}, D_{\text{dom}}} (\mathcal{L} + \lambda \mathcal{L}_d) \end{aligned} \quad (5)$$

where λ is a weighting factor to balance the effect between distance metrics and adversarial components. In this paper, we use cosine similarity as the distance measure:

$$f_d(\mathbf{x}, \mathbf{x}') = 1 - \frac{\mathbf{x}^\top \mathbf{x}'}{\|\mathbf{x}\| \|\mathbf{x}'\|} \quad (6)$$

We train our model adversarially to learn the mappings bidirectionally by encouraging that the resulting pairs be indistinguishable from genuine pairs, and the direction that was generated remain as indiscernible as possible.

Datasets Corpus	Tatoeba								Europarl							
	deu→eng		eng→deu		spa→eng		eng→spa		deu→eng		eng→deu		spa→eng		eng→spa	
Precision@ <i>k</i> (%)	<i>k</i> =1	<i>k</i> =5	<i>k</i> =1	<i>k</i> =5	<i>k</i> =1	<i>k</i> =5	<i>k</i> =1	<i>k</i> =5	<i>k</i> =1	<i>k</i> =5	<i>k</i> =1	<i>k</i> =5	<i>k</i> =1	<i>k</i> =5	<i>k</i> =1	<i>k</i> =5
Mikolov et al. (2013)	13.6	20.9	12.8	21.7	31.1	46.9	24.7	38.3	6.4	13.4	7.8	14.6	13.1	21.6	12.4	22.1
Dinu et al. (2014)	17.3	30.8	22.7	36.4	35.4	52.8	30.5	46.2	13.8	21.9	11.9	21.1	18.5	28.6	15.7	27.3
Smith et al. (2017)	27.6	42.6	26.0	40.0	40.5	55.2	35.0	48.1	13.9	24.3	19.3	27.5	47.2	58.6	39.5	50.8
Lample et al. (2018)	10.0	19.7	11.2	18.7	29.1	40.3	30.5	43.5	20.8	30.2	22.9	35.2	28.2	39.7	30.0	42.7
BERT	10.0	17.3	12.4	21.1	25.6	38.2	24.9	33.8	5.9	8.9	5.8	9.6	22.5	39.9	21.7	36.5
seq2seq NMT	6.5	9.0	7.3	10.1	9.9	12.5	9.9	13.1	5.0	8.7	5.1	6.3	7.0	9.2	7.8	10.3
fairseq NMT	9.3	19.4	15.6	24.4	19.9	29.8	18.5	31.0	12.2	16.1	7.3	11.8	18.4	27.9	17.0	25.2
Conditional GAN	16.0	36.3	18.8	37.5	16.6	34.2	16.6	32.3	9.6	17.1	7.9	15.0	18.6	29.6	17.3	27.6
ABSent (ours)	46.2	65.5	43.8	65.9	38.5	64.3	37.4	60.5	27.5	40.4	26.2	39.2	49.4	62.5	47.7	62.2

Table 1: Results of sentence embedding mapping experiment in terms of precision@1 and precision@5. Our proposed ABSent only utilizes 20% of parallel training data and the equal size of unparallelled data for training, while all baselines take 100% training data with parallel labels. The best-performing method is highlighted in bold.

Dataset Corpus	Tatoeba				Europarl			
	deu→eng	eng→spa	spa→eng	eng→deu	deu→eng	eng→spa	spa→eng	eng→deu
$U \setminus \text{mis}$	25.1	28.0	27.1	24.4	19.8	19.6	38.6	36.7
U	26.9	29.1	27.4	25.3	21.5	20.0	39.3	36.9
$A \setminus \text{mis}$	32.0	32.9	30.7	29.6	23.9	22.8	43.0	44.5
A	34.4	34.9	31.8	30.0	24.7	23.1	45.0	45.8
A_s	27.4	28.1	28.0	26.2	15.6	14.5	24.6	22.8
A_w	24.0	23.7	22.5	21.5	21.9	22.4	37.1	38.4

Table 2: Comparison of precision@1 between ABSent (A), uni-Sent (U), and further modified settings. $A \setminus \text{mis}$, $U \setminus \text{mis}$ are two variants without mismatch loss term. A_s swaps the weighting strategy (TF-IDF weighting for Tatoeba but not for Europarl). A_w stands for first aligning word-level representations and then generating sentence embeddings.

Language	LASER		ABSent (ours)		#test
	#train	Acc (%)	#train	Acc (%)	
Armenian (hye)	6k	32.21	5.3k	28.56	742
Irish (gle)	732	4.20	700	5.23	1k
Kazakh (kaz)	4k	17.39	3.9k	18.14	575

Table 3: On low-resource languages, simple word vector averages obtain comparable results to the richly supervised LASER model, despite lack of word order information.

3.3 Zero-Shot Multilingual Setting

Our model can also be easily extended to align sentences between two languages X_1 and X_2 in a zero-shot manner without any parallel data between them. The zero-shot multilingual task involves jointly projecting two languages X_1 and X_2 to a common target language Y given only limited parallel data and non-aligned data connecting each to Y . As input, we have labeled data $\mathcal{D}_1^{(i)}$ and unlabeled data $\mathcal{D}_{\text{ul}}^{(i)}$ for the language pairs (X_i, Y) ($i = 1, 2$). However, we do not observe any direct relationship between the two source languages X_1 and X_2 in the training data.

In this case, we adopt the same framework as in the previous section except that the generator G_x is expected to learn two language projections from X_1 to Y and from X_2 to Y . Let $\mathcal{L}^{(i)}$ be the loss function defined in Equation 5 for cross-lingual X_i and Y . The overall loss function for multilingual mapping is simply $\mathcal{L}' = \frac{1}{2} (\mathcal{L}^{(1)} + \mathcal{L}^{(2)})$.

3.4 Sentence Representation and Mapping

With regard to obtaining the sentence representations, we adopt pre-trained word vectors (Bojanowski et al. 2016) that are available for numerous languages trained on Wikipedia using fastText. We ensure the same input embeddings in training and evaluation for the baselines as in our model. Based on the results of Wieting et al. (2015) and Arora, Liang, and Ma (2017), we adopt simple (weighted) averages of word vectors, which are surprisingly powerful, although our method could also be applied to other sentence embedding methods.

Given source sentence embeddings $\{\mathbf{x}\}$ and target sentence embeddings $\{\mathbf{y}\}$ acquired as described above, we can train the generators G_X and G_Y through the joint loss function in Equation 5. Subsequently, we evaluate the obtained transformation via a standard sentence retrieval task. For each source sentence embedding \mathbf{x} , we compute its k nearest neighbours in terms of the distance function f_d among all target embeddings. The corresponding k target sentences are regarded as the candidate set of mapping results.

4 Experiments

In this section, we extensively evaluate the effectiveness of our ABSent method compared with state-of-the-art approaches on two heterogeneous real-world corpora.

4.1 Experimental Setup

Datasets. We evaluate the precision of our approach on the Europarl parallel corpus and on extracted from the Tatoeba service¹, which provides translations of commonly used phrases that might be useful to language learners. We focus on German and English as well as Spanish and English translation retrieval. For the English \longleftrightarrow German datasets, we take 160k pairs as the training set and 1,600 pairs as the test set in both datasets. For the English \longleftrightarrow Spanish datasets, we take 60k pairs as training and 600 pairs as test data for the Tatoeba corpus, and 130k as training and 1,300 as test for the Europarl corpus. However, to emphasize that our model can cope with very limited amounts of parallel data, we solely make use of just 20% of the parallel training

¹<http://tatoeba.org>

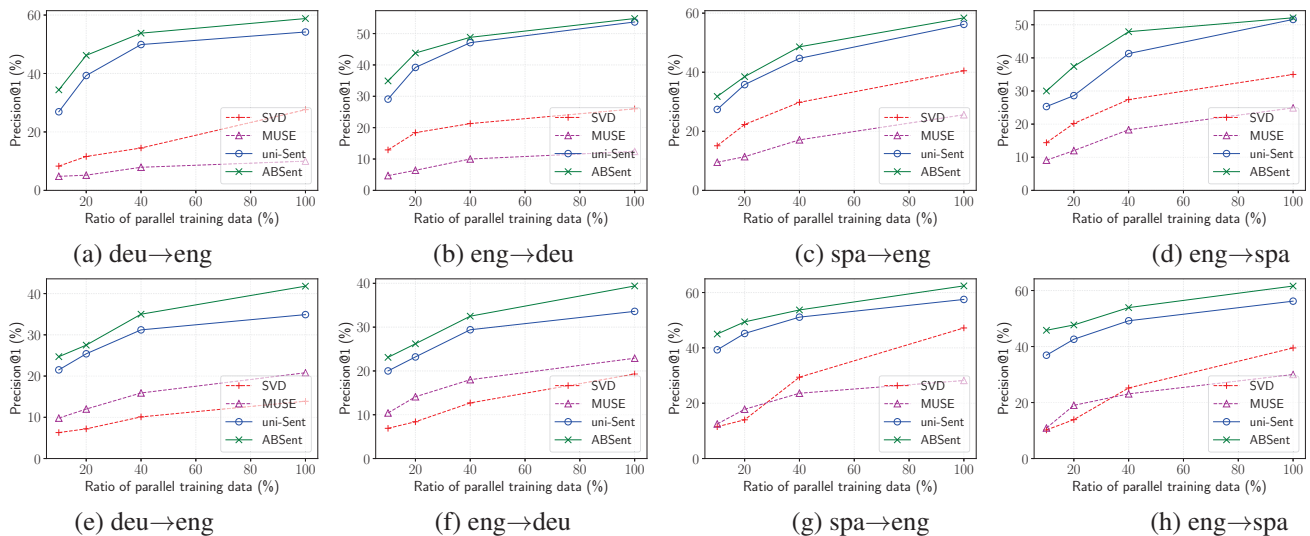


Figure 3: Performance of our proposed ABSent and uni-Sent with two baseline experiments under different ratios of parallel training data on the Tatoeba dataset (Figures (a)-(d)) and Europarl dataset (Figures (e)-(h)).

data when training our model, while all the baseline methods exploit 100% of the parallel training data. Given a set of training pairs, we randomly sample false pairs of the same size as the respective parallel data.

Baselines. For comparison, we consider as baselines the linear transformation methods by (Mikolov, Le, and Sutskever 2013), Dinu, Lazaridou, and Baroni (2014), and Smith et al. (2017), the supervised version of MUSE (Lample et al. 2018) with cross-domain similarity local scaling. We also use the multilingual version of BERT (Devlin et al. 2018), using the standard method for sentence-level representations based on the [CLS] token² to generate sentence vectors that are already multilingual without further projection. This is to assess how far we can take simple word vector averages in comparison to powerful alternatives.

We further consider a seq2seq (Sutskever, Vinyals, and Le 2014) NMT baseline jointly trained to translate language X to Y as well as to monolingually auto-encode sentences from language Y back to X . We use two different encoders with a shared decoder such that the two encoders produce latent representations in the same space. This allows us to save the latent sentence embeddings for evaluation rather than generate an output translation.

Additionally, we consider the fairseq NMT (Gehring et al. 2017) approach based on a convolutional encoder model, which constructs latent representations hierarchically.

Finally, we investigate a Conditional GAN (Mirza and Osindero 2014), for which we use our model but do not consider any fake pairs or mismatch pairs for training.

Parameter Settings. Both generators G_X and G_Y consist of three fully connected layers with hidden sizes of 512,

²As provided by based on bert-as-service: <https://github.com/hanxiao/bert-as-service>

1024, 512, respectively. Each hidden layer is connected with a *BatchNorm* layer and the *ReLU* activation function. The final activation function is *tanh*. Both discriminators D_{real} and D_{dom} take as input two embeddings, followed by three fully connected layers of sizes 512, 1024, 512 with concatenation. Each hidden layer is connected with a *leaky ReLU* activation function (0.2), while the output is activated by a *sigmoid* function. We rely on Adam optimization with an initial learning rate of 0.002 and a batch size of 128.

4.2 Main Results

We assess the quality of the mapping by considering the ranking of the ground truth paired target sentence. The overall quality across all test set instances is given by the precision@ k metric, which, following previous work (Mikolov, Le, and Sutskever 2013) in this area, is defined as the ratio of test set instances for which the correct target is among the top k . Then we repeat the same evaluation process for all four datasets from the two corpora.

The results are reported in Table 1. Recall that we only use 20% of parallel sentences (true pairs) to train our model while all the other baselines utilize 100% of parallel sentence pairs for the training. We observe that our ABSent approach still significantly outperforms other baselines by a large margin. Take the deu→eng data from Tatoeba as an example. Our method achieves a precision@1 of 46.2% and a precision@5 of 65.5%, which are 18.6 and 22.9 *absolute* percentage points higher than the respective results of the best baseline. Similar trends can be observed for other datasets. Note that the results for different languages are not fully comparable due to different sizes of training data.

4.3 Detailed Analysis

Influence of Bi-directional Transformation. We evaluate how the bi-directional mapping strategy affects the effectiveness of our model. Taking German sentence embeddings

Example	German Sentence	English Sentences ranked by similarity score
Tatoeba (high-quality)	Ich finde keine Worte.	(1) I am at a loss for words. (0.852) (2) I just don't know what to say. (0.847) (3) Tom has turned twenty. (0.836)
Europarl (high-quality)	Notwendig ist die Interoperabilität der Set-Top-Boxen.	(1) Set top boxes must be compatible with each other. (0.895) (2) People are showing great solidarity at local and regional level and help is being mobilised at national level. (0.858) (3) I support the proposed deadline for the Commission being 30 September in proposed amendment. (0.846)
Europarl (medium-quality)	Zunächst ist es zu begrüßen, dass der Universaldienst zwar einfache, aber keine breitbandigen Internetanschlüsse umfassen soll.	(1) In the past, the international community has done itself credit by prohibiting anti-personnel mines on these grounds. It should now, by the same token, ban weapons containing depleted uranium. (0.866) (2) This is because it provides clarity and therefore does not expose public services to the attacks which would otherwise have been levelled at them. (0.863) (3) That does not mean, however, that we do not still see much room for improvement, as other speakers have pointed out as well, and some of our wishes have not been fulfilled. (0.854) (4) Firstly, the fact that the universal service is to include simple but not broadband Internet connections is to be welcomed. (0.852)
Europarl (low-quality)	Sie teilt die Auffassung, dass mit dem Universaldienst nicht nur die geographische Abdeckung gewährleistet werden soll.	(1) At the European Council in Gothenburg at the end of this week, the Swedish Presidency will point out the need to discuss these issues within the European Union in order to develop a concrete basis allowing powerful action by the European Union on these vital issues. (0.862) (2) It could have expressed a lot more in the way of hopes for the future. (0.860) (3) So what does this railway package contain? (0.859) ... (6) The Commission shares the view that universal service is not just about getting geographical coverage right. (0.814)

Table 4: Examples of English sentences as dom neighbors for German sentences in Europarl and Tatoeba. (parallel ratio=10%)

Datasets	Tatoeba				Europarl			
	deu→spa		spa→deu		deu→spa		spa→deu	
Metrics	Acc	P@5	Acc	P@5	Acc	P@5	Acc	P@5
(Baseline A)	8.2	11.5	9.0	15.4	7.5	16.4	8.3	16.1
(Baseline B)	12.3	26.1	13.5	27.2	8.1	14.9	9.5	19.5
(Baseline C)	20.8	39.1	17.4	37.2	21.1	34.2	20.9	35.4
(Baseline D)	15.1	27.5	15.6	27.9	20.9	31.6	21.4	33.1
(Baseline E)	19.2	29.4	18.7	30.4	19.3	34.6	20.1	32.4
seq2seq	5.2	6.9	5.6	7.0	4.1	5.9	5.0	6.4
NMT	13.3	26.2	15.1	29.0	22.8	31.8	23.4	27.9
fairseq	13.3	26.2	15.1	29.0	22.8	31.8	23.4	27.9
NMT	13.3	26.2	15.1	29.0	22.8	31.8	23.4	27.9
ABSent	27.3	49.2	26.6	55.4	30.8	45.3	28.6	39.8

Table 5: Results of multilingual sentence embedding mapping experiment in terms of accuracy (%) and precision@5 (%). The best-performing method is highlighted in bold. (Baseline A: Mikolov, Le, and Sutskever (2013), Baseline B: Dinu, Lazaridou, and Baroni (2014), Baseline C: Smith et al. (2017), Baseline D: Lample et al. (2018), Baseline E: Schwenk and Douze (2017))

as source and English sentence embedding as the target, we train our model to map the German sentence embeddings to the corresponding English vector space and align the sentence with the same meaning. We obtain the unidirectional transformation model from German to English, which we refer to as *uni-Sent*. Then we repeat the same process for English to German, Spanish to English, and English to Spanish. Thus, this model only acquires the ability to conduct a uni-

directional transformation between two languages, since the bi-directional discriminator D_{dom} is omitted. In this case, we only learn the generator G_X to map the source to the target domain. The goal is to optimize

$$\begin{aligned}
\mathcal{L}_{\text{real}} &= \mathbb{E}_{\mathbf{x}, \mathbf{y}} [\log(D_{\text{real}}(\mathbf{x}, \mathbf{y}))] \\
&+ \mathbb{E}_{\mathbf{x}} [\log(1 - D_{\text{real}}(\mathbf{x}, G(\mathbf{x})))] \\
G_X^* &= \min_{G_X} \max_{D_{\text{real}}} (\mathcal{L}_{\text{real}} + \mathcal{L}_{\text{mis}} + \lambda \mathcal{L}_d)
\end{aligned} \tag{7}$$

where \mathcal{L}_{mis} and \mathcal{L}_d stay the same as in Equations 3 and 1. The results in Figure 3 shows that the effectiveness of the method is improved substantially by the introduction of bidirectional learning. This demonstrates that the bidirectional training method not only enables a simultaneous transformation between the two language sentence embeddings, but also delivers better results. One possible reason is the effectiveness of D_{dom} , which can regularize the directional uncertainty among two language sentence embedding transformations. A unidirectional transformation does not bear this benefit.

Influence of Ratio of Parallel Training Corpus. Next, we study how various ratios of available parallel data affect the effectiveness of our model. Apart from using 20% of parallel data in the training corpus, we also evaluate using 10%, 40%, and 100% as ratios of parallel training sentence pairs. We randomly sample the mismatch pairs to be of equal size as the respective parallel data. The results are also depicted in Figure 3. We observe that the precision improves as the ratio of parallel labelled sentence increases.

Influence of Mismatch Loss. Recall that our ABSent method incorporates a custom mismatch loss in Equation 3. This experiment aims to study how this loss function influences the effectiveness of our proposed models. For simplicity, we refer to a model X without the mismatch loss as $X \setminus mis$, where X can be either *uni-Sent* or *ABSent*. The results in Table 2 demonstrate the effectiveness of the mismatch strategy of introducing false pairs. We take 10% of parallel training pairs, leaving other parameters as in the main experiments. Generally, bringing in mismatch pairs into the training improves the effectiveness of both our ABSent model and the uni-Sent version.

Influence of Weighting strategy For the input sentence embeddings of ABSent, we take the average of word embeddings and get normalized to obtain sentence representations for the Tatoeba corpus. For the Europarl corpus, we define the sentence embeddings to be the normalized sum over the word vectors multiplied with TF-IDF scores for a weighted average. As a comparison experiment, we swap the weighting strategy in A_s while keeping other parameters as for ABSent, i.e., we impose TF-IDF weighted vectors for Tatoeba, while using vanilla word averages for Europarl.

The effectiveness of A_s is shown in Table 2. Choosing an appropriate weighting strategy boosts the experimental results. Since TF-IDF weights words in accordance with their assumed importance, for the Europarl corpus, the 2σ volume of sentence lengths falls between 35 and 60, while for the Tatoeba corpus it is between 8 and 16. In such circumstances, words that appear more frequently in long sentences, especially function words such as 'a', 'the', etc. ought to have a lower weight, while infrequent ones ought to have a higher weight. For short sentences, it may make sense to even consider words with higher frequency, so as to not neglect their semantic contribution to the sentence. Thus, simple averaging works better than TF-IDF weighted averaging for the Tatoeba corpus.

First Aligning Words. As we can see in Table 2, the accuracy drops quite notably compared to the regular ABSent approach if we first align words and then create the sentence vectors. The only difference between A_w and ABSent is that we first align individual word vectors using our method and then average them (TF-IDF weighted averaging for the Europarl corpus) to generate sentence embeddings in the target vector space. For this comparison, we take 10% of parallel word pairs for training. We conjecture that this approach is less able to account for variation in the meaning of a word across different sentences.

Qualitative Analysis. We additionally provide representative examples of nearest neighbours, showcasing typical high-quality, medium quality, as well as low-quality transformation results for English–German in Table 4. The lower quality results in some cases highlight the limits of average word vector based embeddings, as they disregard word order and may lose semantically salient signals. This problem can

be overcome by applying our method using more semantically sophisticated methods to obtain sentence embeddings. Although many such methods require extensive training and in some cases also rich supervision, the advantage of our model is that we can rely on just limited parallel data to project a resource-poor language into the embedding space of a resource-rich language such as English for which such sentence embedding methods are readily available.

4.4 Mapping of Low-Resource Languages

In order to better evaluate the effectiveness and robustness of our model for diverse languages, we conduct additional experiments on low-resource languages. For comparison, we consider the state-of-the-art massively multilingual neural MT model LASER³ (Artetxe and Schwenk 2018). We evaluate the mapping of low-resource languages with English on the test sets provided by them. However, as they only provide test sets but not training sets, our model is trained on comparably sized training data obtained via random sampling from Tatoeba, OpenSubtitles2018, Global Voices.⁴ We adopt equivalent preprocessing steps such as filtering certain special characters and eliminating duplicate pairs.

The results in Table 3 confirm that simple word averages can be aligned with a broadly similar level of accuracy. This is obtained although our method does not have access to word order information and is not trained on the rich massively multilingual data used to train LASER, but only on the respective single language pairs.

4.5 Zero-Shot Multilingual Mapping

We also evaluate multilingual training, which entails mapping two source languages (Spanish and German) both to the same target language space (English) without any parallel data connecting the two source languages.

From both Tatoeba and Europarl, we each take 44,280 Spanish \longleftrightarrow English sentence pairs and 44,280 German \longleftrightarrow English sentence pairs. To train the baselines, we also collect the same number of German \longleftrightarrow Spanish pairs for which the bilingual baselines make use of dedicated supervision, while our method does not receive any pairings at all for this language pair. Additionally, our proposed method only utilizes 20% of parallel training data and an equal amount of non-parallel data for training, while all the baselines take 100% training data with parallel labels. During the training process, we alternate over mini-batches with Spanish–English pairings and German–English pairings. The number of test queries is 490 for all experiments.

The results are reported in Table 5. Though seq-to-seq models can learn full-fledged neural translation models, they do not fare particularly well in resource-constrained scenarios with limited training data. Particularly, even with a parallel sentence pair percentage of just 20%, our model outperforms many baselines that utilize the total amount of training data. Moreover, retrieval accuracies between two source languages German and Spanish obtained by our model are

³<https://github.com/facebookresearch/LASER>

⁴Available from <http://opus.nlpl.eu/>. For Irish, we incorporate some additional training data from the EUbookshop dataset

very competitive with baselines receiving supervision for that language pair. Note that in our method, we do not provide any direct pairwise mapping. This proves the effectiveness of our zero-shot Multilingual mapping.

5 Conclusion

Our study shows that despite their simplicity, word vector averages can serve as reasonably strong cross-lingually projectable sentence representations. To this end, we have presented the ABSent model to align such representations via an adversarial approach that requires only small amounts of parallel data. We obtain competitive results, although our method does not obtain any information about the word order in the input sentences. Our results in a series of retrieval experiments on both short and long sentences outperform previous work by a substantial margin.

References

- Arora, S.; Liang, Y.; and Ma, T. 2017. A simple but tough-to-beat baseline for sentence embeddings. In *Proc. of ICLR*.
- Artetxe, M., and Schwenk, H. 2018. Massively multilingual sentence embeddings for zero-shot cross-lingual transfer and beyond. *CoRR* abs/1812.10464.
- Artetxe, M.; Labaka, G.; and Agirre, E. 2017. Learning bilingual word embeddings with (almost) no bilingual data. In *Proceedings of ACL*.
- Bojanowski, P.; Grave, E.; Joulin, A.; and Mikolov, T. 2016. Enriching word vectors with subword information. *arXiv 1607.04606*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Dinu, G.; Lazaridou, A.; and Baroni, M. 2014. Improving zero-shot learning by mitigating the hubness problem. *arXiv 1412.6568*.
- Dong, X., and De Melo, G. 2018. Cross-lingual propagation for deep sentiment analysis. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Faruqui, M., and Dyer, C. 2014. Improving vector space word representations using multilingual correlation. In *Proceedings of EACL 2014*.
- Ferrero, J.; Agnes, F.; Besacier, L.; and Schwab, D. 2017. Using word embedding for cross-language plagiarism detection. In *Proceedings of ACL 2017*.
- Gan, Z.; Chen, L.; Wang, W.; Pu, Y.; Zhang, Y.; Liu, H.; Li, C.; and Carin, L. 2017. Triangle generative adversarial networks. In *Proceedings of NIPS*.
- Gehring, J.; Auli, M.; Grangier, D.; Yarats, D.; and Dauphin, Y. 2017. Convolutional Sequence to Sequence Learning. *arXiv 1705.03122*.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. In *Proceedings of NIPS*, 2672–2680.
- Isola, P.; Zhu, J.-Y.; Zhou, T.; and Efros, A. A. 2017. Image-to-image translation with conditional adversarial networks. In *Proceedings of CVPR*.
- Kiros, R.; Zhu, Y.; Salakhutdinov, R.; Zemel, R. S.; Torralba, A.; Urtasun, R.; and Fidler, S. 2015. Skip-thought vectors. In *Proceedings of NIPS*. MIT Press.
- Lample, G.; Conneau, A.; Ranzato, M.; Denoyer, L.; and Jégou, H. 2018. Word translation without parallel data. In *In Proceeding of ICML 2018*.
- Le, Q., and Mikolov, T. 2014. Distributed representations of sentences and documents. In *Proceedings of ICML 2014*. PMLR.
- Liu, B.; Zhu, Y.; Fu, Z.; de Melo, G.; and Elgammal, A. 2019a. Oogan: Disentangling gan with one-hot sampling and orthogonal regularization. *arXiv:1905.10836*.
- Liu, J.; Lin, Y.; Liu, Z.; and Sun, M. 2019b. Xqa: A cross-lingual open-domain question answering dataset. In *Proceedings ACL 2019*.
- Mikolov, T.; Chen, K.; Corrado, G.; and Dean, J. 2013. Efficient estimation of word representations in vector space. *arXiv 1301.3781*.
- Mikolov, T.; Le, Q. V.; and Sutskever, I. 2013. Exploiting similarities among languages for machine translation. *arXiv 1309.4168*.
- Mirza, M., and Osindero, S. 2014. Conditional generative adversarial nets. *CoRR* abs/1411.1784.
- Nakashole, N., and Flauger, R. 2018. Characterizing departures from linearity in word translation. In *ACL*.
- Ruder, S. 2017. A survey of cross-lingual embedding models. *arXiv 1706.04902*.
- Schwenk, H., and Douze, M. 2017. Learning joint multilingual sentence representations with neural machine translation. In *Rep4NLP@ACL*.
- Smith, S. L.; Turban, D. H. P.; Hamblin, S.; and Hammerla, N. Y. 2017. Offline bilingual word vectors, orthogonal transformations and the inverted softmax. *arXiv 1702.03859*.
- Subramanian, S.; Trischler, A.; Bengio, Y.; and Pal, C. J. 2018. Learning general purpose distributed sentence representations via large scale multi-task learning. *arXiv 1804.00079*.
- Sutskever, I.; Vinyals, O.; and Le, Q. V. 2014. Sequence to sequence learning with neural networks. In *Proceedings of NIPS*, 3104–3112.
- Wang, X.; Pham, H. Q.; Arthur, P.; and Neubig, G. 2019. Multilingual neural machine translation with soft decoupled encoding. In *Proc. of ICLR*.
- Wieting, J.; Bansal, M.; Gimpel, K.; and Livescu, K. 2015. Towards universal paraphrastic sentence embeddings. *arXiv 1511.08198*.
- Xian, Y.; Fu, Z.; Muthukrishnan, S.; de Melo, G.; and Zhang, Y. 2019. Reinforcement knowledge graph reasoning for explainable recommendation. In *Proceedings of SIGIR 2019*.
- Xing, C.; Wang, D.; Liu, C. C.; and Lin, Y. 2015. Normalized word embedding and orthogonal transform for bilingual word translation. In *HLT-NAACL*.
- Yang, Y.; Ábrego, G. H.; Yuan, S.; Guo, M.; Shen, Q.; Cer, D.; Sung, Y.-H.; Strope, B.; and Kurzweil, R. 2019. Improving multilingual sentence embedding using bi-directional dual encoder with additive margin softmax. *CoRR* abs/1902.08564.
- Zhang, M.; Liu, Y.; Luan, H.; and Sun, M. 2017. Adversarial training for unsupervised bilingual lexicon induction. In *Proceedings of ACL*.
- Zhu, J.-Y.; Park, T.; Isola, P.; and Efros, A. A. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. *arXiv 1703.10593*.
- Zou, W. Y.; Socher, R.; Cer, D.; and Manning, C. D. 2013. Bilingual word embeddings for phrase-based machine translation. In *Proceedings of EMNLP*, 1393–1398.