

Bilingual Dictionary Based Neural Machine Translation without Using Parallel Sentences

Xiangyu Duan¹, Baijun Ji¹, Hao Jia¹, Min Tan¹, Min Zhang^{1*},
Boxing Chen², Weihua Luo², Yue Zhang³

¹ Institute of Artificial Intelligence, School of Computer Science and Technology,
Soochow university

² Alibaba DAMO Academy

³ School of Engineering, Westlake University

{xiangyuduan,minzhang}@suda.edu.cn; {bjji,hjia,mtan2017}@stu.suda.edu.cn;

{boxing.cbx,weihua.luowh}@alibaba-inc.com; yue.zhang@wias.org.cn

Abstract

In this paper, we propose a new task of machine translation (MT), which is based on no parallel sentences but can refer to a ground-truth bilingual dictionary. Motivated by the ability of a monolingual speaker learning to translate via looking up the bilingual dictionary, we propose the task to see how much potential an MT system can attain using the bilingual dictionary and large scale monolingual corpora, while is independent on parallel sentences. We propose anchored training (AT) to tackle the task. AT uses the bilingual dictionary to establish anchoring points for closing the gap between source language and target language. Experiments on various language pairs show that our approaches are significantly better than various baselines, including dictionary-based word-by-word translation, dictionary-supervised cross-lingual word embedding transformation, and unsupervised MT. On distant language pairs that are hard for unsupervised MT to perform well, AT performs remarkably better, achieving performances comparable to supervised SMT trained on more than 4M parallel sentences¹.

1 Introduction

Motivated by a monolingual speaker acquiring translation ability by referring to a bilingual dictionary, we propose a novel MT task that no parallel sentences are available, while a ground-truth bilingual dictionary and large-scale monolingual corpora can be utilized. This task departs from unsupervised MT task that no parallel resources, including the ground-truth bilingual dictionary, are allowed to utilize (Artetxe et al., 2018c; Lample et al., 2018b). This task is also distinct to

supervised/semi-supervised MT task that mainly depends on parallel sentences (Bahdanau et al., 2015; Gehring et al., 2017; Vaswani et al., 2017; Chen et al., 2018; Sennrich et al., 2016a).

The bilingual dictionary is often utilized as a seed in bilingual lexicon induction (BLI) that aims to induce more word pairs within the language pair (Mikolov et al., 2013). Another utilization of the bilingual dictionary is for translating low-frequency words in supervised NMT (Arthur et al., 2016; Zhang and Zong, 2016). We are the first to utilize the bilingual dictionary and the large scale monolingual corpora to see how much potential an MT system can achieve without using parallel sentences. This is different from using artificial bilingual dictionaries generated by unsupervised BLI for initializing an unsupervised MT system (Artetxe et al., 2018c,b; Lample et al., 2018a), we use the ground-truth bilingual dictionary and apply it throughout the training process.

We propose Anchored Training (AT) to tackle this task. Since word representations are learned over monolingual corpora without any parallel sentence supervision, the representation distances between source language and target language are often quite large, leading to significant translation difficulty. As one solution, AT selects words covered by the bilingual dictionary as anchoring points to drive the distance between the source language space and the target language space closer so that translation between the two languages becomes easier. Furthermore, we propose Bi-view AT that places anchors based on either source language view or target language view, and combines both views to enhance the translation quality.

Experiments on various language pairs show that AT performs significantly better than various baselines, including word-by-word translation through looking up the dictionary, unsupervised MT, and dictionary-supervised cross-lingual word embed-

* Corresponding Author.

¹Code is available at <https://github.com/mttravel/Dictionary-based-MT>

ding transformation to make distances between both languages closer. Bi-view AT further improves AT performance due to mutual strengthening of both views of the monolingual data. When combined with cross-lingual pretraining (Lample and Conneau, 2019), Bi-view AT achieves performances comparable to traditional SMT systems trained on more than 4M parallel sentences. The main contributions of this paper are as follows:

- A novel MT task is proposed which can only use the ground-truth bilingual dictionary and monolingual corpora, while is independent on parallel sentences.
- AT is proposed as a solution to the task. AT uses the bilingual dictionary to place anchors that can encourage monolingual spaces of both languages to become closer so that translation becomes easier.
- The detailed evaluation on various language pairs shows that AT, especially Bi-view AT, performs significantly better than various methods, including word-by-word translation, unsupervised MT, and cross-lingual embedding transformation. On distant language pairs that unsupervised MT struggled to be effective, AT and Bi-view AT perform remarkably better.

2 Related Work

The bilingual dictionaries used in previous works are mainly for bilingual lexicon induction (BLI), which independently learns the embedding in each language using monolingual corpora, and then learns a transformation from one embedding space to another by minimizing squared euclidean distances between all word pairs in the dictionary (Mikolov et al., 2013; Artetxe et al., 2016). Later efforts for BLI include optimizing the transformation further through new training objectives, constraints, or normalizations (Xing et al., 2015; Lazaridou et al., 2015; Zhang et al., 2016; Artetxe et al., 2016; Smith et al., 2017; Faruqui and Dyer, 2014; Lu et al., 2015). Besides, the bilingual dictionary is also used for supervised NMT which requires large-scale parallel sentences (Arthur et al., 2016; Zhang and Zong, 2016). To our knowledge, we are the first to use the bilingual dictionary for MT without using any parallel sentences.

Our work is closely related to unsupervised NMT (UNMT) (Artetxe et al., 2018c; Lample

et al., 2018b; Yang et al., 2018; Sun et al., 2019), which does not use parallel sentences neither. The difference is that UNMT may use the artificial dictionary generated by unsupervised BLI for initialization (Artetxe et al., 2018c; Lample et al., 2018a) or abandon the artificial dictionary by using joint BPE so that multiple BPE units can be shared by both languages (Lample et al., 2018b). We use the ground-truth dictionary instead and apply it throughout a novel training process. UNMT works well on close language pairs such as English-French, while performs remarkably bad on distant language pairs in which aligning the embeddings of both side languages is quite challenging. We use the ground-truth dictionary to alleviate such problem, and experiments on distant language pairs show the necessity of using the bilingual dictionary.

Other utilizations of the bilingual dictionary for tasks beyond MT include cross-lingual dependency parsing (Xiao and Guo, 2014), unsupervised cross-lingual part-of-speech tagging and semi-supervised cross-lingual super sense tagging (Gouws and Søgaard, 2015), multilingual word embedding training (Ammar et al., 2016; Duong et al., 2016), and transfer learning for low-resource language modeling (Cohn et al., 2017).

3 Our Approach

There are multiple freely available bilingual dictionaries such as Muse dictionary² (Conneau et al., 2018), Wiktionary³, and PanLex⁴. We adopt Muse dictionary which contains 110 large-scale ground-truth bilingual dictionaries.

We propose to inject the bilingual dictionary into the MT training by placing anchoring points on the large scale monolingual corpora to drive the semantic spaces of both languages becoming closer so that MT training without parallel sentences becomes easier. We present the proposed Anchored Training (AT) and Bi-view AT in the following.

3.1 Anchored Training (AT)

Since word embeddings are trained on monolingual corpora independently, the embedding spaces of both languages are quite different, leading to significant translation difficulty. AT forces words of a translation pair to share the same word embedding as an anchor. We place multiple anchors by

²<https://github.com/facebookresearch/MUSE>

³https://en.wiktionary.org/wiki/Wiktionary:Main_Page

⁴<https://panlex.org/>

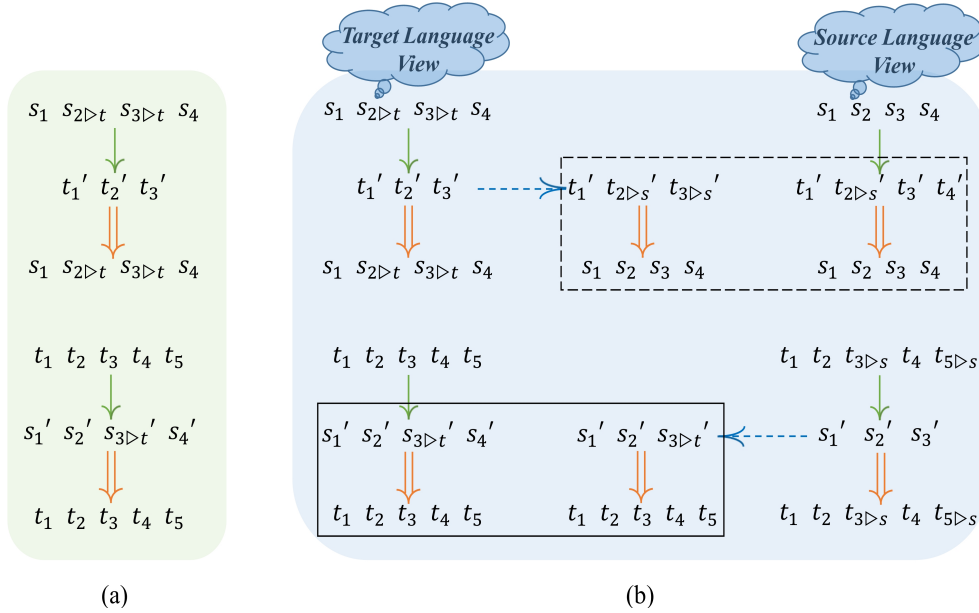


Figure 1: Illustration of (a) AT and (b) Bi-view AT. We use a source language sentence “ $s_1 s_2 s_3 s_4$ ” and a target language sentence “ $t_1 t_2 t_3 t_4 t_5$ ” from the large-scale monolingual corpora as an example. \triangleright denotes an anchoring point which replaces a word with its translation based on the bilingual dictionary. Thin arrows of \downarrow denote NMT decoding, thick arrows of \Downarrow denote training an NMT model, $-\rightarrow$ and $-\leftarrow$ denote generating the anchored sentence based on the dictionary. Words with primes such as t_1' denote the decoding output of a thin arrow.

selecting words covered by the bilingual dictionary. With stable anchors, the embedding spaces of both languages become more and more close during the AT process.

As illustrated in Figure 1 (a), given the source sentence “ $s_1 s_2 s_3 s_4$ ” with words of s_2 and s_3 being covered by the bilingual dictionary, we replace the two words with their translation words according to the dictionary. This results in the source sentence “ $s_1 s_2 \triangleright t s_3 \triangleright t s_4$ ”, of which $s_2 \triangleright t$ and $s_3 \triangleright t$ serve as the **anchors** which are actually the target language words obtained by translating s_2 and s_3 according to the dictionary, respectively. Through the anchors, some words on the source side share the same word embeddings with the corresponding words on the target side. The AT process will strengthen the consistency of embedding spaces of both languages based on these anchors.

The training process illustrated in Figure 1 (a) consists of a mutual back-translation procedure. The anchored source sentence “ $s_1 s_2 \triangleright t s_3 \triangleright t s_4$ ” is translated into target sentence “ $t_1' t_2' t_3'$ ” by using source-to-target **decoding**, then “ $t_1' t_2' t_3'$ ” and “ $s_1 s_2 \triangleright t s_3 \triangleright t s_4$ ” constitute a sentence pair for **training** the target-to-source translation model. In contrast, the target sentence “ $t_1 t_2 t_3 t_4 t_5$ ” is translated into anchored source sentence “ $s_1' s_2' s_3 \triangleright t' s_4'$ ” by using target-to-source decoding, then both sen-

tences constitute a sentence pair for training the source-to-target translation model. Note that during **training** the translation model, the input sentences are always pseudo sentences generated by decoding an MT model, while the output sentences are always true or anchored true sentences. Beside this mutual back-translation procedure, a denoising procedure used in unsupervised MT (Lample et al., 2018b) is also adopted. The deletion and permutation noises are added to the source/target sentence, and the translation model is also trained to denoise them into the original source/target sentence.

During testing, a source sentence is transformed into an anchored sentence at first by looking up the bilingual dictionary. Then we use the source-to-target model trained in the AT process to decode the anchored sentence.

We use Transformer architecture (Vaswani et al., 2017) as our translation model with four stacked layers in both encoder and decoder. In the encoder, we force the last three layers shared by both languages, and leave the first layer not shared. In the decoder, we force the first three layers shared by both languages, and leave the last layer not shared. Such architecture is designed to capture both common and specific characteristics of the two languages in one model for the training.

3.2 Bi-view AT

AT as illustrated in Figure 1 (a) actually tries to model the sentences of both languages in the target language view with partial source words replaced with the target words and the full target language sentence. Bi-view AT enhances AT by adding another language view. Figure 1 (b) adds the source language view shown in the right part to accompany with the target language view of Figure 1 (a). In particular, the target language sentence “ $t_1 t_2 t_3 t_4 t_5$ ” is in the form of “ $t_1 t_2 t_{3 \triangleright s} t_4 t_{5 \triangleright s}$ ” after looking up the bilingual dictionary. Such partial target words replaced with the source words and the full source language sentence “ $s_1 s_2 s_3 s_4$ ” constitute the source language view.

Based on the target language view shown in the left part and the source language view shown in the right part, we further combine both views through the pseudo sentences denoted by primes in Figure 1 (b). As shown by “ \dashrightarrow ” in Figure 1 (b), “ $t_1' t_2' t_3'$ ” is further transformed into “ $t_1' t_{2 \triangleright s}' t_{3 \triangleright s}'$ ” by looking up the bilingual dictionary. Similarly, “ $s_1' s_2' s_3'$ ” is further transformed into “ $s_1' s_{2 \triangleright t}' s_{3 \triangleright t}'$ ” as shown by “ \dashleftarrow ”. Finally, solid line box represents training the source-to-target model on data from both views, and dashed line box represents training the target-to-source model on data from both views.

Bi-view AT starts from training both views in parallel. After both views converge, we generate pseudo sentences in both the solid line box and the dashed line box, and pair these pseudo sentences (as input) with genuine sentences (as output) to train the corresponding translation model. This generation and training process iterates until Bi-view AT converges. Through such rich views, the translation models of both directions are mutually strengthened.

3.3 Anchored Cross-lingual Pretraining (ACP)

Cross-lingual pretraining has demonstrated effectiveness on tasks such as cross-lingual classification, unsupervised MT (Lample and Conneau, 2019). It is conducted over large monolingual corpora by masking random words and training to predict them as a cloze task. Instead, we propose ACP to pretrain on data that is obtained by transforming the genuine monolingual corpora of both languages into the anchored version. For example, words in the source language corpus that are cov-

ered by the bilingual dictionary are replaced with their translation words respectively. Such words are anchoring points that can drive the pretraining to close the gap between the source language space and the target language space better than the original pretraining method of Lample and Conneau (2019) does as evidenced by the experiments in section 4.5. Such anchored source language corpus and the genuine target language corpus constitute the target language view for ACP.

ACP can be conducted in either the source language view or the target language view. After ACP, each of them is used to initialize the encoder of the corresponding AT system.

3.4 Training Procedure

For AT, the pseudo sentence generation step and NMT training step are interleaved. Take the target language view AT shown in Figure 1 (a) for example, we extract anchored source sentences as one batch, and decode them into pseudo target sentences; then we use the same batch to train the NMT model of target-to-anchored source. In the meantime, a batch of target sentences are decoded into pseudo anchored source sentences, and then we use the same batch to train the NMT model of anchored source-to-target. The above process repeats until AT converges.

For Bi-view AT, after each mono-view AT converging, we set larger batch for generating pseudo sentences as shown in solid/dashed line boxes in Figure 1 (b), and train the corresponding NMT model using the same batch.

For ACP, we follow XLM procedure (Lample and Conneau, 2019), and conduct pretraining on the anchored monolingual corpora concatenated with the genuine corpora of the other language.

4 Experimentation

We conduct experiments on English-French, English-Russian, and English-Chinese translation to check the potential of our MT system with only bilingual dictionary and large scale monolingual corpora. The English-French task deals with the translation between close-related languages, while the English-Russian and English-Chinese tasks deal with the translation between distant languages that do not share the same alphabets.

4.1 Datasets

For English-French translation task, we use the monolingual data released by XLM (Lample and Conneau, 2019)⁵. For English-Russian translation task, we use the monolingual data identical to Lample et al.(2018a), which uses all available sentences for the WMT monolingual News Crawl datasets from years 2007 to 2017. For English-Chinese translation task, we extract Chinese sentences from half of the 4.4M parallel sentences from LDC, and extract English sentences from the complementary half. We use WMT *newstest-2013/2014*, WMT *newstest-2015/2016*, and NIST2006/NIST2002 as validation/test sets for English-French, English-Russian, and English-Chinese, respectively.

For cross-lingual pretraining, we extract raw sentences from Wikipedia dumps, which contain 80M, 60M, 13M, 5.5M monolingual sentences for English, French, Russian, and Chinese, respectively.

Muse ground-truth bilingual dictionaries are used for our dictionary-related experiments. If a word has multiple translations, we select the translation word that appears most frequently in the monolingual corpus. Table 1 summarizes the number of word pairs and their coverage on the monolingual corpora on the source side.

	entry no.	coverage
fr→en	97,046	60.91%
en→fr	94,719	69.77%
ru→en	45,065	65.43%
en→ru	42,725	88.77%
zh→en	13,749	50.20%
en→zh	32,495	47.02%

Table 1: Statistics of Muse bilingual dictionaries.

4.2 Experiment Settings

For AT/Bi-view AT without cross-lingual pretraining, we use Transformer with 4 layers, 512 embedding/hidden units, and 2048 feed-forward filter size, for fair comparison to UNMT (Lample et al., 2018b). For AT/Bi-view AT with ACP, we set Transformer with 6 layers, 1024 embedding/hidden units, and 4096 feed-forward filter size for a fair comparison to XLM (Lample and Conneau, 2019).

We conduct joint byte-pair encoding (BPE) on the monolingual corpora of both languages with a shared vocabulary of 60k tokens for both English-French and English-Russian tasks, and 40k tokens for English-Chinese task (Sennrich et al., 2016b).

⁵<https://github.com/facebookresearch/XLM/blob/master/get-data-nmt.sh>

During training, we set the batch size to 32 and limit the sentence length to 100 BPE tokens. We employ the Adam optimizer with $lr = 0.0001$, $t_{warm_up} = 4000$ and $dropout = 0.1$. At decoding time, we generate greedily with length penalty $\alpha = 1.0$.

4.3 Baselines

- Word-by-word translation by looking up the ground truth dictionary or the artificial dictionary generated by Conneau et al. (2018).
- Unsupervised NMT (UNMT) that does not rely on any parallel resources (Lample et al., 2018b)⁶. Besides, cross-lingual pretraining (XLM) based UNMT (Lample and Conneau, 2019)⁷, is also set as a stronger baseline (XLM+UNMT).
- We implement a UNMT initialized by Unsupervised Word Embedding Transformation (UNMT+UWET) as a baseline(Artetxe et al., 2018d). The transformation function is learned in an unsupervised way without using any ground-truth bilingual dictionaries (Conneau et al., 2018)⁸.
- We also implement a UNMT system initialized by Supervised Word Embedding Transformation (UNMT+SWET) as a baseline. Instead of UWET used in Artetxe et al. (2018d), we use the ground-truth bilingual dictionary as the supervision signal to train the transformation function for transforming the source word embeddings into the target language space (Conneau et al., 2018). After such initialization, the gap between the embedding spaces of both languages is narrowed for easy UNMT training.

4.4 Experimental Results: without Cross-lingual Pretraining

The upper part of Table 2 presents the results of various baselines and our AT approaches. AT and Bi-view AT significantly outperform the baselines, and Bi-view AT is consistently better than AT. Detailed comparisons are listed as below:

Results of Word-by-word Translation

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

⁷<https://github.com/facebookresearch/XLM>

⁸<https://github.com/facebookresearch/MUSE>

system	fr → en	en → fr	ru → en	en → ru	zh → en	en → zh
Without Cross-lingual Pre-training						
Word-by-word using artificial dictionary	7.76	4.88	3.05	1.60	1.99	1.14
Word-by-word using ground-truth dictionary	7.97	6.61	4.17	2.81	2.68	1.79
UNMT (Lample et al., 2018b)	24.02	25.10	9.09	7.98	1.50	0.45
UNMT+SWET	21.11	21.22	9.79	4.07	19.78	7.84
UNMT+UWET	19.80	21.27	8.79	6.21	15.54	6.62
AT	25.07	26.36	10.20	9.91	19.83	9.18
Bi-view AT	27.11	27.54	12.85	10.64	21.16	11.23
With Cross-lingual Pre-training						
XLM+UNMT (Lample and Conneau, 2019)	33.28	35.10	17.39	13.29	20.68	11.28
ACP+AT	33.51	36.15	16.41	15.43	26.80	13.91
ACP+Bi-view AT	34.05	36.56	20.09	17.62	30.12	17.05
Supervised SMT	-	-	21.48	14.54	31.86	16.55

Table 2: Experiment results evaluated by BLEU using the multi-bleu script.

It shows that using the ground-truth dictionary is slightly better than using the artificial one generated by Conneau et al. (2018). Both performances are remarkably bad, indicating that simple word-by-word translation is not qualified as an MT method. More effective utilization of the bilingual dictionary is needed to improve the translation performance.

Comparison between UNMT and UNMT with WET Initialization

UNMT-related systems generally improves the performance of the word-by-word translation. On the close-related language pair of English-French, UNMT is better than UNMT+UWET/SWET. This is partly because there are numerous BPE units shared by both English and French, enabling easy establishing the shared word embedding space of both languages. In contrast, WET that transforms the source word embedding into the target language space seems not a necessary initialization step since shared BPE units already establish the shared space.

On distant language pairs, UNMT does not have an advantage over UNMT with WET initialization. Especially on English-Chinese, UNMT performs extremely bad, even worse than the word-by-word translation method. We argue that this is because the BPE units shared by both languages are so few that UNMT fails to align the language spaces. In contrast, using the bilingual dictionary greatly alleviate such problem for distant language pairs. UNMT+SWET, which transforms the source word embedding into the target word embedding space supervised by the bilingual dictionary, outperforms UNMT by more than 18 BLEU points on Chinese-to-English and more than 7 BLEU points

on English-to-Chinese. This indicates the necessity of the bilingual dictionary for translation between distant language pairs.

Comparison between AT/Bi-view AT and The Baselines

Our proposed AT approaches significantly outperform the baselines. The baselines of using the ground-truth bilingual dictionary, i.e., word-by-word translation using the dictionary and UNMT+SWET that uses the dictionary to supervise the word embedding transformation, are inferior to our AT approaches.

The AT approaches consistently improves the performances over both close-related language pair of English-French and distant language pairs of English-Russian and English-Chinese. Our Bi-view AT achieves the best performance on all language pairs.

4.5 Experimental Results: with Cross-lingual Pretraining

The bottom part of Table 2 reports performances of UNMT with XLM, which conducts the cross-lingual pretraining on concatenated non-parallel corpora (Lample and Conneau, 2019), and performances of our AT/Bi-view AT with the anchored cross-lingual pretraining, i.e., ACP. The results show that our proposed AT approaches are still superior when equipped with the cross-lingual pretraining.

UNMT obtains great improvement when combined with XLM, achieving state-of-the-art unsupervised MT performance better than Unsupervised SMT (Artetxe et al., 2019) and Unsupervised NMT (Lample et al., 2018b) across close and distant language pairs.




Figure 2: Visualization of the bilingual word embeddings after Bi-view AT.

ACP+AT/Bi-view AT performs consistently superior to XLM+UNMT. Especially on distant language pairs, ACP+Bi-view AT gains 2.7-9.4 BLEU improvements over the strong XLM+UNMT. This indicates that AT/Bi-view AT with ACP builds closer language spaces via anchored pretraining and anchored training. We present such advantage in the analyses of Section 4.6.

Comparison with Supervised SMT

To check the ability of our system using only the dictionary and non-parallel corpora, we make the comparison to supervised SMT trained on over 4M parallel sentences, which are from WMT19 for English-Russian and from LDC for English-Chinese. We use Moses⁹ as the supervised SMT system with a 5-gram language model trained on the target language part of the parallel corpora.

The bottom part of Table 2 shows that ACP+Bi-view AT performs comparable to supervised SMT, and performs even better on English-to-Russian and English-to-Chinese.

4.6 Analyses

We analyze the cross-lingual property of our approaches in both word level and sentence level. We also compare the performances between the ground-truth dictionary and the artificial dictionary. In the end, we vary the size of the bilingual dictionary and report its impact on the AT training.

⁹<http://www.statmt.org/ Moses/>. We use the default setting of Moses.

Effect on Bilingual Word Embeddings

As shown in Figure 2, we depict the word embeddings of some sampled words in English-Chinese after our Bi-view AT. The dimensions of the embedding vectors are reduced to two by using T-SNE and are visualized by the visualization tool in Tensorflow¹⁰.

We sample the English words that are not covered by the dictionary at first, then search their nearest Chinese neighbors in the embedding space. It shows that the words which constitute a new ground-truth translation pair do appear as neighboring points in the 2-dimensional visualization of Figure 2.

Precision of New Word Pairs

We go on with studying bilingual word embedding by quantitative analysis of the new word pairs, which are detected by searching bilingual words that are neighbors in the word embedding space, and evaluate them using the ground-truth bilingual dictionary. In particular, we split the Muse dictionary of Chinese-to-English into standard training set and test set as in BLI (Artetxe et al., 2018a). The training set is used for the dictionary-based systems, including our AT/Bi-view AT, UNMT+SWET, and Muse, which is a BLI toolkit. The test set is used to evaluate these systems by computing the precision of discovered translation words given the source words in the test set. The neighborhood is computed by CSLS distance (Conneau et al., 2018).

Table 3 shows the precision, where precision@k indicates the accuracy of top-k predicted candidate. Muse induces new word pairs through either the supervised way or the unsupervised way. MuseSupervised is better than MuseUnsupervised since it is supervised by the ground-truth bilingual dictionary. Our AT/Bi-view AT surpasses MuseSupervised by a large margin. UNMT+SWET/UWET also obtains good performance through the word embedding transformation. Bi-view AT significantly surpasses UNMT+SWET/UWET in precision@5 and precision@10, while is worse than them in precision@1. This indicates that Bi-view AT can produce better *n*-best translation words that are beneficial for NMT beam decoding to find better translations.

Through the word level analysis, we can see that AT/Bi-view AT leads to more consistent word em-

¹⁰<https://projector.tensorflow.org/>

	MuseUnsupervised	MuseSupervised	UNMT+SWET	UNMT+UWET	AT	Bi-view AT
Precision@1	30.51	35.38	48.01	45.85	43.32	45.49
Precision@5	55.42	58.48	68.05	67.15	68.23	72.02
Precision@10	62.45	63.18	72.02	72.20	73.83	76.71

Table 3: Precision of Discovered New Word Pairs.

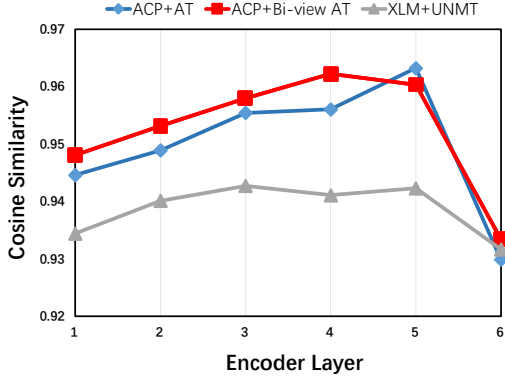


Figure 3: Sentence level cosine similarity of the parallel sentences on each encoder layer.

bedding space shared by both languages, making the translation between both languages easier.

Sentence Level Similarity of Parallel Sentences

We check the sentence level representational invariance across languages for the cross-lingual pretraining methods. In detail, following Arivazhagan et al. (2018), we adopt max-pooling operation to collect the sentence representation of each encoder layer for all Chinese-to-English sentence pairs in the test set. Then we calculate the cosine similarity for each sentence pair and average all cosine scores.

Figure 3 shows the sentence level cosine similarity. ACP+Bi-view AT consistently has a higher similarity for parallel sentences than XLM+UNMT on all encoder layers. When compare Bi-view AT and AT, the Bi-view AT is better on more encoder layers.

We can see that in both word level and sentence level analysis, our AT methods achieve better cross-lingual invariance, significantly reduce the gap between the source language space and the target language space, leading to decreased translation difficulty between both languages.

Ground-Truth Dictionary Vs Artificial Dictionary

Table 4 presents the comparison in English-Chinese. The ground-truth dictionary is from the Muse dictionary deposit, and the artificial dictio-

	Ground-Truth Dict.		Artificial Dict.	
	zh→en	en→zh	zh→en	en→zh
AT	19.83	9.18	16.7	6.98
Bi-view AT	21.16	11.23	18.23	8.50

Table 4: BLEU of AT methods using either the ground-truth dictionary or the artificial dictionary.

XLM+UNMT	20.68
ACP+AT with 1/4 of the dictionary	22.84
ACP+AT with 1/2 of the dictionary	24.32
ACP+AT with the full dictionary	26.80

Table 5: BLEU of ACP+AT using different size of the dictionary in zh→en translation.

nary is generated by unsupervised BLI (Conneau et al., 2018). We extract top- n word pairs as the artificial dictionary, where n is the same as the number of entries in the ground-truth dictionary.

Both dictionaries use AT methods for translation. As shown in Table 4, the ground-truth dictionary performs significantly better than the artificial dictionary in both methods and both translation directions.

The Effect of The Dictionary Size

We randomly select a portion of the ground-truth bilingual dictionary to study the effect of the dictionary size on the performance. Table 5 reports the performances of ACP+AT using a quarter or a half of the zh→en dictionary.

It shows that, in comparison to the baseline of XLM+UNMT that does not use a dictionary, a quarter of the dictionary consisting of around 3k word pairs is capable of improving the performance significantly. More word pairs in the dictionary lead to better translation results, suggesting that expanding the size of the current Muse dictionary via collecting various dictionaries built by human experts may improve the translation performance further.

5 Discussion and Future Work

In the literature of unsupervised MT that only uses non-parallel corpora, Unsupervised SMT (USMT) and Unsupervised NMT (UNMT) are

complementary to each other. Combining them (USMT+UNMT) achieves significant improvement over the individual system, and performs comparable to XLM+UNMT (Lample et al., 2018b; Artetxe et al., 2019).

We have set XLM+UNMT as a stronger baseline, and our ACP+AT/Bi-view AT surpasses it significantly. By referring to the literature of unsupervised MT, we can opt to combine ACP+AT/Bi-view AT with SMT. We leave it as a future work.

6 Conclusion

In this paper, we explore how much potential an MT system can achieve when only using a bilingual dictionary and large-scale monolingual corpora. This task simulates people acquiring translation ability via looking up the dictionary and depending on no parallel sentence examples. We propose to tackle the task by injecting the bilingual dictionary into MT via anchored training that drives both language spaces closer so that the translation becomes easier. Experiments show that, on both close language pairs and distant language pairs, our proposed approach effectively reduces the gap between the source language space and the target language space, leading to significant improvement of translation quality over the MT approaches that do not use the dictionary and the approaches that use the dictionary to supervise the cross-lingual word embedding transformation.

Acknowledgments

The authors would like to thank the anonymous reviewers for the helpful comments. This work was supported by National Natural Science Foundation of China (Grant No. 61525205, 61673289), National Key R&D Program of China (Grant No. 2016YFE0132100), and was also partially supported by the joint research project of Alibaba and Soochow University.

References

Waleed Ammar, George Mulcaire, Yulia Tsvetkov, Guillaume Lample, Chris Dyer, and Noah A. Smith. 2016. Massively multilingual word embeddings. *ArXiv:1602.01925*.

Naveen Arivazhagan, Ankur Bapna, Orhan Firat, Roei Aharoni, Melvin Johnson, and Wolfgang Macherey. 2018. The missing ingredient in zero-shot neural machine translation. *ArXiv:1903.07091*.

Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2016. Learning principled bilingual mappings of word embeddings while preserving monolingual invariance. In *Proc. of EMNLP*.

Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2018a. A robust self-learning method for fully unsupervised cross-lingual mappings of word embeddings. In *Proc. of ACL*.

Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2018b. Unsupervised statistical machine translation. In *Proc. of EMNLP*.

Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2019. An effective approach to unsupervised machine translation. In *Proc. of ACL*.

Mikel Artetxe, Gorka Labaka, Eneko Agirre, and Kyunghyun Cho. 2018c. Unsupervised neural machine translation. In *Proc. of ICLR*.

Mikel Artetxe, Gorka Labaka, Eneko Agirre, and Kyunghyun Cho. 2018d. Unsupervised neural machine translation. In *Proc. of ICLR*.

Philip Arthur, Graham Neubig, and Satoshi Nakamura. 2016. Incorporating discrete translation lexicons into neural machine translation. In *Proc. of EMNLP*.

Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *Proc. of ICLR*.

Mia Xu Chen, Orhan Firat, Ankur Bapna, Melvin Johnson, Wolfgang Macherey, George Foster, Llion Jones, Niki Parmar, Michael Schuster, Zhi-Feng Chen, Yonghui Wu, and Macduff Hughes. 2018. The best of both worlds: Combining recent advances in neural machine translation. In *Proc. of ACL*.

Trevor Cohn, Steven Bird, Graham Neubig, Oliver Adams, and Adam J. Makarucha. 2017. Cross-lingual word embeddings for low-resource language modeling. In *Proc. of EACL*.

Alexis Conneau, Guillaume Lample, Marc’Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. 2018. Word translation without parallel data. In *Proc. of ICLR*.

Long Duong, Hiroshi Kanayama, Tengfei Ma, Steven Bird, and Trevor Cohn. 2016. Learning crosslingual word embeddings without bilingual corpora. In *Proc. of EMNLP*.

Manaal Faruqui and Chris Dyer. 2014. Improving vector space word representations using multilingual correlation. In *Proc. of EACL*.

Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann Dauphin. 2017. Convolutional sequence to sequence learning. In *Proc. of ICML*.

Stephan Gouws and Anders Søgaard. 2015. Simple task-specific bilingual word embeddings. In *Proc. of HLT-NAACL*.

- Guillaume Lample and Alexis Conneau. 2019. Cross-lingual language model pretraining. *ArXiv:1901.07291*.
- Guillaume Lample, Ludovic Denoyer, and Marc’Aurelio Ranzato. 2018a. Unsupervised machine translation using monolingual corpora only. In *Proc. of ICLR*.
- Guillaume Lample, Myle Ott, Alexis Conneau, Ludovic Denoyer, and Marc’Aurelio Ranzato. 2018b. Phrase-based & neural unsupervised machine translation. In *Proc. of EMNLP*.
- Angeliki Lazaridou, Georgiana Dinu, and Marco Baroni. 2015. Hubness and pollution: Delving into cross-space mapping for zero-shot learning. In *Proc. of ACL*.
- Ang Lu, Weiran Wang, Mohit Bansal, Kevin Gimpel, and Karen Livescu. 2015. Deep multilingual correlation for improved word embeddings. In *Proc. of HLT-NAACL*.
- Tomas Mikolov, Quoc V. Le, and Ilya Sutskever. 2013. Exploiting similarities among languages for machine translation. *ArXiv:1309.4168*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016a. Improving neural machine translation models with monolingual data. In *Proc. of ACL*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016b. Neural machine translation of rare words with subword units. In *Proc. of ACL*.
- 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 *Proc. of ICLR*.
- Haipeng Sun, Rui Wang, Kehai Chen, Masao Utiyama, Eiichiro Sumita, and Tiejun Zhao. 2019. Unsupervised bilingual word embedding agreement for unsupervised neural machine translation. In *Proc. of ACL*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proc. of NIPS*.
- Min Xiao and Yuhong Guo. 2014. Distributed word representation learning for cross-lingual dependency parsing. In *Proc. of CoNLL*.
- Chao Xing, Dong Wang, Chao Liu, and Yiye Lin. 2015. Normalized word embedding and orthogonal transform for bilingual word translation. In *Proc. of HLT-NAACL*.
- Zhen Yang, Wei Chen, Feng Wang, and Bo Xu. 2018. Unsupervised neural machine translation with weight sharing. In *Proc. of ACL*.
- Jiajun Zhang and Chengqing Zong. 2016. Bridging neural machine translation and bilingual dictionaries. *ArXiv:1610.07272*.
- Yuan Zhang, David Gaddy, Regina Barzilay, and Tommi S. Jaakkola. 2016. Ten pairs to tag - multilingual pos tagging via coarse mapping between embeddings. In *Proc. of HLT-NAACL*.